

New Avenues in Opinion Mining and Sentiment Analysis

Erik Cambria, *National University of Singapore*

Björn Schuller, *Technical University of Munich*

Yunqing Xia, *Tsinghua University*

Catherine Havasi, *Massachusetts Institute of Technology*

Others' opinions can be crucial when it's time to make a decision or choose among multiple options. When those choices involve valuable resources (for example, spending time and money to buy products or services) people often rely on their peers' past experiences. Until recently, the main sources

of information were friends and specialized magazine or websites. Now, the "social web" provides new tools to efficiently create and share ideas with everyone connected to the World Wide Web. Forums, blogs, social networks, and content-sharing services help people share useful information. This information is unstructured, however, and because it's produced for human consumption, it's not something that's "machine processable." Capturing public opinion about social events, political movements, company strategies, marketing campaigns, and product preferences is garnering increasing interest from the scientific community (for the exciting open challenges), and from the business world (for the remarkable marketing fallouts and for possible financial market prediction). The resulting emerging fields are *opinion mining* and *sentiment analysis*. Although commonly used interchangeably to denote the same field of study, opinion mining

and sentiment analysis actually focus on polarity detection and emotion recognition, respectively. Because the identification of sentiment is often exploited for detecting polarity, however, the two fields are usually combined under the same umbrella or even used as synonyms. Both fields use data mining and natural language processing (NLP) techniques to discover, retrieve, and distill information and opinions from the World Wide Web's vast textual information.

Mining opinions and sentiments from natural language is challenging, because it requires a deep understanding of the explicit and implicit, regular and irregular, and syntactical and semantic language rules. Sentiment analysis researchers struggle with NLP's unresolved problems: coreference resolution, negation handling, anaphora resolution, named-entity recognition, and word-sense disambiguation. Opinion mining is a very restricted NLP problem,

The Web holds valuable, vast, and unstructured information about public opinion. Here, the history, current use, and future of opinion mining and sentiment analysis are discussed, along with relevant techniques and tools.

because the system only needs to understand the positive or negative sentiments of each sentence and the target entities or topics. Therefore, sentiment analysis is an opportunity for NLP researchers to make tangible progress on all fronts of NLP, and potentially have a huge practical impact.

Many companies use opinion mining and sentiment analysis as part of their research. For instance, companies use opinion mining to create and automatically maintain review and opinion-aggregation websites. Their systems continuously gather a wide array of information from the Web, such as product reviews, brand perception, and political issues. Other systems might also use opinion mining and sentiment analysis as subcomponent technology to improve customer relationship management and recommendation systems through positive and negative customer feedback. Similarly, opinion mining and sentiment analysis might detect and exclude “flames” (overly heated or antagonistic language) in social communication and enhance antispy systems.

Companies use sentiment analysis to develop marketing strategies by assessing and predicting public attitudes toward their brand. Research and development focuses on designing automatic tools that crawl online reviews and condense the information gathered. Numerous companies already provide tools that track public viewpoints on a large scale by offering graphical summarizations of trends and opinions in the blogosphere. Developing opinion-tracking systems is commercially important.

Also, several tools already exist to help companies extract and analyze information from blogs about large-scale trends in customers’ opinions about products; those tools include

SenticNet (<http://sentic.net>), Luminoso (<http://luminoso.com>), Factiva (<http://dowjones.com/factiva>), Attensity (<http://attensity.com>), and Converseon (<http://converseon.com>). Most existing tools and research, however, are limited to polarity evaluation or mood classification according to a limited set of emotions. Such methods mainly rely on parts of text in which people explicitly express emotional states, and therefore the tools can’t capture a reviewer’s implicitly expressed opinion or sentiment. To better consider the state of this field, we discuss here the past, present, and future trends of sentiment analysis by delving into the evolution of opinion mining systems. More comprehensive surveys on sentiment analysis can be found elsewhere.¹⁻³

Common Sentiment Analysis Tasks

The basic task of opinion mining is *polarity classification*. Polarity classification occurs when a piece of text stating an opinion on a single issue is classified as one of two opposing sentiments. Reviews such as “thumbs up” versus “thumbs down,” or “like” versus “dislike” are examples of polarity classification. Polarity classifications also identify pro and con expressions in online reviews and help make the product evaluations more credible.

Agreement detection is another form of binary sentiment classification. Agreement detection determines whether a pair of text documents should receive the same or different sentiment-related labels. After the system identifies the polarity classification, it might assign *degrees of positivity* to the polarity—that is, it might locate the opinion on a continuum between positive and negative. Also, it can classify multimedia resources according to mood and

emotional content for purposes such as affective human-machine interaction, troll filtering, and cyber-issue detection. If the text doesn’t contain strong opinions or covers more than one issue or item, new challenges arise, such as subjectivity detection and opinion-target identification. Distinguishing between subjective and objective text helps classify the sentiment. Moreover, a piece of text might 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, a system performs sentiment analysis over on-topic documents—using, for example, the results of a topic-based search engine. However, several studies suggest that managing these two tasks jointly might benefit overall performance. For example, a document’s off-topic passages might contain irrelevant affective information and create inaccurate global-sentiment polarity about the main topic. Also, a document might contain information on multiple topics that interest the user. In such instances, it’s important to identify topics and separate the opinions associated with each topic.

Evolution of Opinion Mining

Currently, opinion mining and sentiment analysis rely on vector extraction to represent the most salient and important text features. We can use this vector to classify the most relevant features. Two commonly used features are *term frequency* and *presence*.

Presence is a binary-valued feature vector in which the entries indicate only whether a term occurs (value 1) or doesn’t (value 0). Presence forms a more effective basis to review polarity classification and reveals an interesting difference: although recurrent keywords indicate a topic, repeated

terms might not reflect the overall sentiment.

It's possible to add other term-based features to the features vector. *Position* refers to how a token's position in a text unit might affect the text's sentiment. Further, we might consider *presence n-grams*—typically bigrams and trigrams—to be useful features. Some methods also rely on the distance between terms. General textual analysis uses part of speech (POS) information (for example, nouns, adjectives, adverbs, and verbs) as a basic form of word-sense disambiguation. Certain adjectives are good indicators of sentiment and guide feature selection to classify the sentiment. Also, selected phrases chosen by pre-specified POS patterns, usually including an adjective or adverb, help detect sentiments.

Some researchers have developed other text mapping techniques that assign labels to predefined categories or real numbers representing the degree of polarity. These approaches are strictly bound by domain and topic. Moreover, most research on sentiment analysis focuses on text written in English and, consequently, most of the resources developed (such as sentiment lexicons and corpora) are in English. Applying this research to other languages is a domain adaptation problem.

From Heuristics to Discourse Structure

In some unsupervised learning approaches, a *sentiment lexicon* is generated and later used to determine the text unit's degree of positivity or subjectivity. Creating the sentiment lexicon through unsupervised polarity or subjectivity labeling of words or phrases is crucial.¹ The sentiment lexicon identifies a term or a phrase's prior polarity or prior subjectivity, which in turn helps identify contextual polarity or subjectivity. Early works focused

mostly on linguistic heuristics. For example, in their work on polarity classification, Vasileios Hatzivassiloglou and Kathleen McKeown discuss how two classes of interest represent opposites.⁴ These opposite constraints help the system with label decisions.

These approaches were unable to detect novel expression of sentiment. Consequently, later work focused on propagating the valence of seed words (for which the polarity is known) to terms that co-occur with them in general text (or in dictionary glosses) or to synonyms and words that co-occur with them in other WordNet-defined relations. For example, Ana-Maria Popescu and Oren Etzioni proposed an *iterative collective labeling* algorithm.⁵ This algorithm starts with a global word label computed over a large collection of generic topic text. Gradually the algorithm redefines the label with more specificity: first to a specific review corpus, then specific to a product feature, and finally to a label specific to the context in which the word occurs. Benjamin Snyder and Regina Barzilay similarly explored using discourse information to infer relationships between product attributes.⁶ They designed a linear classifier that would predict whether all aspects of a product would be given the same rating. Then they combined the prediction with individual-aspect classifiers, which would minimize loss function.

For opinionated documents, such as product reviews, regression techniques are often used to predict the degree of positivity of opinions. Regression techniques implicitly model similar relationships between classes that correspond to points on a scale, such as the number of stars that a reviewer gives.¹ Modeling discourse structure, such as twists and turns in a document, leads to more effective sentiment labeling. In earlier

research, Bo Pang and Lillian Lee attempted to partially address this problem by incorporating location information into the feature set.⁷

More recent studies emphasize the importance of position in sentiment summarization. For example, the incipits of articles in topic-based summarization usually indicate the text's sentiment. However, the last *n* sentences of a product review often best summarize the document's overall sentiment—almost as well as the *n* (automatically computed) of most subjective sentences.⁷ Mahesh Joshi and Carolyn Penstein-Rosé, for example, explored how to use features based on syntactic dependency relations to improve opinion-mining performance.⁸ They converted a transformation of dependency-relation triples into *composite back-off features* that generalize better than the regular, lexicon-based, dependency-relation features.

From Coarse- to Fine-Grained Analysis

We see opinion mining and sentiment analysis research evolving in both technique sophistication and analysis depth. Early on, Bo Pang and her colleagues classified entire documents by overall positive or negative polarity, and also by rating scores of reviews.^{9,10} These documents were mainly supervised, manually labeled samples, such as movie or product reviews explicitly indicating an overall positive or negative opinion.

Opinions and sentiments don't occur only at the document level, nor are they limited to a single valence or target. One document might contain positive and negative opinions toward one or more topics. Hence, later work adopted a *segment-level* opinion analysis that used graph-based techniques to distinguish sentimental from unsentimental sections. Pang and Lee used segment-level opinion

KNOWLEDGE-BASED APPROACHES TO CONCEPT-LEVEL SENTIMENT ANALYSIS

analysis in their work to segment sections of a document by subjectivity. In another study, Peter Turney classified items based on fixed, syntactic phrases used for expressing opinions.¹¹ Finally, Jaap Kamps and his colleagues classified items by *bootstrapping*—using a small set of seed opinion words and a knowledge base such as WordNet.¹²

In another work, Ellen Riloff and Janyce Wiebe reduced text-analysis granularity to the sentence level by using the presence of opinion-bearing lexical items (single words or *n*-grams) to detect subjective sentences.¹³ Soo-Min Kim and Eduard Hovy, instead, used semantic frames that identified sentimental topics (or targets).¹⁴ Reviewers tend to adhere to being either subjective or objective, and that creates continuity among adjacent sentences. Hence, other researchers collectively classify documents by assigning preferences for pairs of nearby sentences.¹⁰

Even sentence-level approaches often fail to discover sentiments about an entity and/or its aspects. To correct that, other researchers adopted an aspect-level approach, wherein an opinion consists of targets and the sentiments associated with them.^{15–17} For example, the sentence “the new iPhone 5’s screen size is amazing, but its battery life is short” evaluates two aspects (opinion targets): the screen size and battery life of the same entity. The sentiment about the iPhone 5’s screen size is positive, but the sentiment about its battery life is negative. Based on this level of analysis, we can produce a structured opinion summary about an entity and its aspects, and can draw more accurate statistics about those aspects.

From Keywords to Concepts

We can study the evolution of sentiment analysis research by the analytical

tokens, or building blocks, and the implicit information associated with those tokens. We can group the existing approaches into four main categories: keyword spotting, lexical affinity, statistical methods, and concept-based techniques.

Keyword spotting. Although the most naïve approach, keyword spotting’s accessibility and economy make it popular. This approach classifies text by affect categories based on the presence of unambiguous affect words such as happy, sad, afraid, and bored. For example, Clark Elliott’s Affective Reasoner watches for 198 affect keywords (such as distressed or enraged), affect intensity modifiers (such as extremely, somewhat, or mildly), and a handful of cue phrases (such as did that and wanted to).¹⁸ Other popular sources of affect words are Andrew Ortony and his colleagues’ Affective Lexicon,¹⁹ which groups terms into affective categories, and Janyce Wiebe and her colleagues’ linguistic annotation scheme.²⁰

Keyword spotting is weak in two areas: it can’t reliably recognize affect-negated words, and it relies on surface features. Although keyword spotting can correctly classify the sentence “today was a happy day” as being affectively positive, it is likely to assign the same classification to a sentence like “today wasn’t a happy day at all.” Also, keyword spotting relies on the presence of obvious affect words that are only surface features of the prose. Sometimes, a sentence conveys affect through underlying meaning rather than affect adjectives. For example, the text “My husband just filed for divorce and he wants to take custody of my children away from me” evokes strong emotions, but uses no affect keywords, and therefore is ineffective. Lexical affinity is slightly more sophisticated than keyword spotting.

Lexical affinity. This approach not only detects obvious affect words, it also assigns arbitrary words a probable “affinity” to particular emotions. For example, lexical affinity might assign the word “accident” a 75-percent probability of indicating a negative affect, as in “car accident” or “hurt by accident.” This approach usually trains probability from linguistic corpora.^{21–23} Although it often outperforms pure keyword spotting, there are two main problems with this approach. First, negated sentences (I avoided an accident) and sentences with other meanings (I met my girlfriend by accident) trick lexical affinity, because they operate solely on the word level. Second, lexical affinity probabilities are often biased toward text of a particular genre, dictated by the linguistic corpora’s source. This makes it difficult to develop a reusable, domain-independent model.

Statistical methods. This approach, which includes Bayesian inference and support vector machines, is popular for affect text classification. Researchers use statistical methods on projects such as Pang’s movie review classifier and many others.^{9,10,15,24} By feeding a machine-learning algorithm a large training corpus of affectively annotated texts, the system might not only learn the affective valence of affect keywords (as in the keyword-spotting approach), but also take into account the valence of other arbitrary keywords (similar to lexical affinity), punctuation, and word co-occurrence frequencies.

Generally, statistical methods are semantically weak, which means that individually—with the exception of obvious affect keywords—a statistical model’s other lexical or co-occurrence elements have little predictive value. As a result, statistical text classifiers only work well when they

receive sufficiently large text input. So, while these methods might be able to affectively classify a user's text on the page level or paragraph level, they don't work well on smaller text units such as sentences or clauses.

Concept-based approaches. These methods use Web ontologies or semantic networks to accomplish semantic text analysis.²⁵⁻²⁷ 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 meaning/features associated with natural language concepts. Superior to purely syntactical techniques, concept-based approaches can detect subtly expressed sentiments. Concept-based approaches can analyze multi-word expressions that don't explicitly convey emotion, but are related to concepts that do.

The concept-based approach relies heavily on the depth and breadth of the knowledge bases it uses. Without a comprehensive resource that encompasses human knowledge, an opinion-mining system will have difficulty grasping the semantics of natural language text. Moreover, the typicality of knowledge bases—that is, the fact that they contain only typical information associated with concepts—limits their capability to handle semantic nuances. Their fixed/flat representation, finally, places bounds on inferences of semantic and affective features associated with concepts.

Multimodal Sentiment Analysis

New sources of opinion mining and sentiment analysis abound. Webcams installed in smartphones, touchpads, or other devices let users post opinions

in an audio or audiovisual format rather than in text. For a rough idea of the amount of material, consider that YouTube users upload two days' worth of video material to its website every minute. Aside from converting spoken language to written text for analysis, the audiovisual format provides an opportunity to mine opinions and sentiment. Many new areas might be useful in opinion mining, such as facial expression, body movement, or a video blogger's choice of music or color filters.

Affect analysis, a related field, addresses the use of linguistic, acoustic, and (potentially) video information. This field focuses on a broader set of emotions or the estimation of continuous emotion primitives; for example, valence can be related to sentiment. In one study, researchers provide recent surveys on spoken and written-language-based analysis; in another study, researchers explore further multimodal combinations.^{28,29} There's almost no research that focuses on multimodal sentiment and opinion analysis. Stephan Raaijmakers and his colleagues fuse acoustic and linguistic information, but that information is based on the transcript of the spoken content rather than on automatic speech recognition output.³⁰ In addition to this research, Louis-Philippe Morency and his colleagues combine acoustic, textual, and video features to assess opinion polarity in 47 YouTube videos.³¹ They demonstrate significant improvement in leave-one-video-out evaluation using Hidden Markov Models for classification. The authors identified polarized words, smiles, gazes, pauses, and voice pitch as relevant features. Again, the researchers relied on transcripts to analyze the text and not the actual spoken word.

Multimodal sentiment analysis hasn't been fully explored, but holds great promise as an application. For example,

it might be extremely valuable when a textual transcript is unavailable, and we need a performance point of view for synergy effects and fail-safeness. In the latter respect, it will be particularly interesting to see further modalities involved—such as physiological and brain signals, along with the use of contextual knowledge. We'll then need to investigate analyses of robustness against disturbances in individual (or all) modalities alongside audiovisual confidence estimation.

Discussion

Gradually, sentiment analysis research is distinguishing itself as a separate field, falling between NLP and natural language understanding. Unlike standard syntactical NLP tasks, such as summarization and auto-categorization, opinion mining mainly focuses on semantic inferences and affective information associated with natural language, and doesn't require a deep understanding of text. We envision sentiment analysis research moving toward content-, concept-, and context-based analysis of natural language text, supported by time-efficient parsing techniques suitable for big social data analysis.³²

Collecting opinions on the Web will still require processing at the content/syntactic level, filtering out unopinionated user-generated content (subjectivity detection) and evaluating the trustworthiness of the opinion and its source. By contrast, concept/semantic analysis infers semantic and affective information associated with natural language opinions, and hence, enables a comparative fine-grained feature-based sentiment analysis. Rather than gathering isolated opinions about a whole item, users generally prefer to compare specific features of different products (for example, the iPhone 5 versus the Galaxy S3 touchscreen) or even sub-features (comparing the

fragility of iPhone 5 and Galaxy S3 touchscreens). To make these comparisons, researchers must construct comprehensive common-knowledge bases to spot features and commonsense bases to detect polarity.³³ Such commonsense bases, in particular, will be key in properly deconstructing natural language text into sentiments—for example, in appraising the concept “small room” as negative for a hotel review and “small queue” as positive for a post office, or the concept “go read the book” as positive for a book review but negative for a movie review.

Context-/intent-level analysis ensures the relevance of the opinions gathered. Social context will continue to gain importance, and an intelligent system will have access to the comprehensive personal information of vast numbers of people. Opinion mining will be specific to each user’s or group of users’ preferences and needs. Opinions won’t be generic, but will reflect their source (for example, a relevant circle of friends or users with similar interests, or the selection of a camera for trekking rather than for night shooting).

The Web has changed from “read-only” to “read-write.” This evolution created enthusiastic users interacting and sharing through social networks, online communities, blogs, wikis, and other collaborative media. Collective knowledge has spread throughout the Web, particularly in areas related to everyday life, such as commerce, tourism, education, and health. Despite significant progress, however, opinion mining and sentiment analysis 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 combine the long tradition of emotion research with their

own discourse, models, and methods. Opinion mining and sentiment analysis are inextricably bound to the affective sciences that attempt to understand human emotions. Affect-sensitive systems and psychological emotion research must develop together.

Recent approaches aim to better grasp the conceptual rules that govern sentiment, as well as the clues that can convert these concepts from realization to verbalization in the human mind. Future opinion-mining systems need broader and deeper common and commonsense knowledge bases. More complete knowledge must be combined with reasoning methods that are more deeply inspired by human thought and psychology. This will lead to a better understanding of natural language opinions and will more efficiently bridge the gap between (unstructured) multimodal information and (structured) machine-processable data.

Blending scientific theories of emotion with the practical engineering goals of analyzing sentiments in natural-language text will lead to more bio-inspired approaches to the design of intelligent opinion-mining systems capable of handling semantic knowledge, making analogies, learning new affective knowledge, and detecting, perceiving, and “feeling” emotions. ■

References

1. 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.
2. B. Liu, *Sentiment Analysis and Opinion Mining*, Morgan and Claypool, 2012.
3. E. Cambria and A. Hussain, *Sentic Computing: Techniques, Tools, and Applications*, Springer, 2012.
4. V. Hatzivassiloglou and K. McKeown, “Predicting the Semantic Orientation of Adjectives,” *Proc. 8th Conf. Assoc. Computational Linguistics European Chapter*, 1997.
5. A. Popescu and O. Etzioni, “Extracting Product Features and Opinions from Reviews,” *Proc. Human Language Technology Conf./Conf. Empirical Methods in Natural Language Processing*, Assoc. for Computational Linguistics, 2005, pp. 339–346.
6. B. Snyder and R. Barzilay, “Multiple Aspect Ranking Using the Good Grief Algorithm,” *Proc. Ann. Conf. North Am. Chapter of the Assoc. for Computational Linguistics*, Assoc. for Computational Linguistics, 2007, article no. N07-1038.
7. B. Pang and L. Lee, “A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts,” *Proc. 42nd Ann. Meeting of the Assoc. for Computational Linguistics*, Assoc. for Computational Linguistics, 2004, pp. 271–278.
8. M. Joshi and C. Penstein-Rosé, “Generalizing Dependency Features for Opinion Mining,” *Proc. 47th Ann. Meeting of the Assoc. for Computational Linguistics and the 4th Int’l Joint Conf. Natural Language Processing of the Asian Federation of Natural Language Processing*, Assoc. for Computational Linguistics, 2009, pp. 313–316.
9. B. Pang, L. Lee, and S. Vaithyanathan, “Thumbs Up? Sentiment Classification Using Machine Learning Techniques,” *Proc. Ann. Conf. Empirical Methods in Natural Language Processing*, Assoc. for Computational Linguistics, 2002, pp. 79–86.
10. B. Pang and L. Lee, “Seeing Stars: Exploiting Class Relationships for Sentiment Categorization with Respect to Rating Scales,” *Proc. 43rd Ann. Assoc. for Computational Linguistics*, Assoc. for Computational Linguistics, 2005, pp. 115–124.
11. P. Turney, “Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews,” *Proc. 40th Ann. Assoc. for Computational Linguistics*, Assoc. for Computational Linguistics, 2002, pp. 417–424.
12. J. Kamps et al., “Using WordNet to Measure Semantic Orientation of Adjectives,” *Proc. 4th Ann. Int’l. Conf.*

THE AUTHORS

Erik Cambria is a research scientist in the Cognitive Science Programme, Temasek Laboratories, National University of Singapore. His research interests include AI, the Semantic Web, natural language processing, and big social data analysis. Cambria has a PhD in computing science and mathematics from the University of Stirling. He is on the editorial board of Springer's *Cognitive Computation* and is the chair of many international conferences. Contact him at cambria@nus.edu.sg.

Björn Schuller leads the Machine Intelligence and Signal Processing group at the Institute for Human-Machine Communication at the Technical University of Munich. His research interests include machine learning, affective computing, and automatic speech recognition. Schuller has a PhD in electrical engineering and information technology from the Technical University of Munich. Contact him at schuller@tum.de.

Yunqing Xia is an associate professor in the Tsinghua National Laboratory for Information Science and Technology at Tsinghua University. His research interests include natural language processing, information retrieval, and text mining. Xia has a PhD in computer science from the Chinese Academy of Science. Contact him at yqxia@tsinghua.edu.cn.

Catherine Havasi is a cofounder of the Open Mind Common Sense project at the Massachusetts Institute of Technology (MIT) Media Lab, where she works as a postdoctoral associate. Her research interests include commonsense reasoning, dimensionality reduction, machine learning, language acquisition, cognitive modeling, and intelligent user interfaces. Havasi has a PhD in computer science from Brandeis University. Contact her at havasi@media.mit.edu.

- Language Resources and Evaluation*, European Language Resources Assoc., 2004, pp. 1115–1118.
13. E. Riloff and J. Wiebe, "Learning Extraction Patterns for Subjective Expressions," *Proc. 2003 Conf. Empirical Methods in Natural Language Processing*, Assoc. for Computational Linguistics, 2003, pp. 105–112.
 14. S. Kim and E. Hovy, "Extracting Opinions, Opinion Holders, and Topics Expressed in Online News Media Text," *Proc. Workshop on Sentiment and Subjectivity in Text*, Assoc. for Computational Linguistics, 2006, pp. 1–8.
 15. M. Hu and B. Liu, "Mining and Summarizing Customer Reviews," *Proc. 10th ACM SIGKDD Conf. Knowledge Discovery and Data Mining*, ACM, 2004, pp. 168–177.
 16. B. Lu et al., "Multi-Aspect Sentiment Analysis with Topic Models," *Proc. Sentiment Elicitation from Natural Text for Information Retrieval and Extraction*, IEEE CS, 2011, pp. 81–88.
 17. G. Di Fabbriozio, A. Aker, and R. Gaizauskas, "STARLET: Multi-Document Summarization of Service and Product Reviews with Balanced Rating Distributions," *Proc. Sentiment Elicitation from Natural Text for Information Retrieval and Extraction*, IEEE, 2011, pp. 67–74.
 18. C.D. Elliott, "The Affective Reasoner: A Process Model of Emotions in a Multi-Agent System," doctoral dissertation, Institute for the Learning Sciences, Northwestern University, 1992.
 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. R. Stevenson, J. Mikels, and T. James, "Characterization of the Affective Norms for English Words by Discrete Emotional Categories," *Behavior Research Methods*, vol. 39, no. 4, 2007, pp. 1020–1024.
 22. S. Somasundaran, J. Wiebe, and J. Ruppenhofer, "Discourse Level Opinion Interpretation," *Proc. 22nd Int'l Conf. Computational Linguistics*, Assoc. for Computational Linguistics, 2008, pp. 801–808.
 23. D. Rao and D. Ravichandran, "Semi-Supervised Polarity Lexicon Induction," *Proc. 12th Conf. European Chapter of the Assoc. for Computational Linguistics*, Assoc. for Computational Linguistics, 2009, pp. 675–682.
 24. L. Nguyen et al., "Predicting Collective Sentiment Dynamics from Time-Series Social Media," *Proc. 18th ACM SIGKDD Conf. Knowledge Discovery and Data Mining*, ACM, 2012, article no. 6.
 25. M. Grassi et al., "Sentic Web: A New Paradigm for Managing Social Media Affective Information," *Cognitive Computation*, vol. 3, no. 3, 2011, pp. 480–489.
 26. D. Olsher, "Full Spectrum Opinion Mining: Integrating Domain, Syntactic and Lexical Knowledge," *Sentiment Elicitation from Natural Text for Information Retrieval and Extraction*, IEEE, 2012, pp. 693–700.
 27. E. Cambria, T. Mazzocco, and A. Husain, "Application of Multi-Dimensional Scaling and Artificial Neural Networks for Biologically Inspired Opinion Mining," *Biologically Inspired Cognitive Architectures*, vol. 4, 2013, pp. 41–53.
 28. B. Schuller et al., "Recognizing Realistic Emotions and Affect in Speech: State of the Art and Lessons Learnt from the First Challenge," *Speech Comm.*, vol. 53, nos. 9/10, 2011, pp. 1062–1087.
 29. 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. 120–135.
 30. S. Raaijmakers, K. Truong, and T. Wilson, "Multimodal Subjectivity Analysis of Multiparty Conversation," *Proc. Conf. Empirical Methods in Natural Language Processing*, Assoc. for Computational Linguistics, 2008, pp. 466–474.
 31. L.-P. Morency, R. Mihalcea, and P. Doshi, "Towards Multimodal Sentiment Analysis: Harvesting Opinions from the Web," *Proc. 13th Int'l Conf. Multimodal Interfaces*, ACM, 2011, pp. 169–176.
 32. E. Cambria et al., "Big Social Data Analysis," *Big Data Computing*, ch. 13, Chapman and Hall/CRC, 2013.
 33. E. Cambria et al., "Semantic Multi-Dimensional Scaling for Open-Domain Sentiment Analysis," *IEEE Intelligent Systems*, preprint, 2012; doi:10.1109/MIS.2012.118.