Word from the Guest Editors

Vivi Nastase
School of Information Technology and Engineering, University of Ottawa
Ottawa, Ontario, Canada

Stan Szpakowicz
School of Information Technology and Engineering, University of Ottawa
Ottawa, Ontario, Canada
Institute of Computer Science, Polish Academy of Sciences
Warsaw, Poland

1 The background

This special issue presents the expanded and improved versions of selected papers from a workshop on Formal and Informal Information Exchange in Negotiations. The workshop took place on May 26-27, 2005, at the University of Ottawa. It was designed to bring together researchers who work on various aspects of interaction in negotiations and those who work in Natural Language Processing or Machine Learning, on problems that might be of interest to negotiation specialists. Such problems include sentiment analysis. Recognizing the sentiments that negotiators express in language (for example, in messages exchanged during electronic negotiations) could offer insight into the negotiation process and a glimpse of the changes in feelings throughout the course of the negotiation. While external factors may influence users at any time, it is likely that they will react fairly consistently to the tone of the partner's messages and offers.

Analysis of texts to recognize attitudes, sentiments and opinions has been receiving more and more attention in recent years. Hearst (1992) proposed direction-based text interpretation to complement topic analysis. Such interpretation showed where a text lies on the axis from *opposed* to *in favour*. Later work on sentiment analysis abandoned the continuous axis in favour of a binary view: recognize positive and negative (or favourable and unfavourable) attitudes. Much research has an economic motivation, as a result of companies wanting to monitor the consumers' satisfaction with their products (Cognitrend[1], Tenorio Research[2]). Opinion recognition may also increase the usefulness of search engines by allowing the user to compare the number and content of positive and negative opinions about a variety of products and services (Das and Chen, 2001), (Dave et al., 2003), (Hu and Liu, 2004), or as part of recommender systems which collect, analyse and summarize user opinions (Tarveen et al., 1997), (Tatemura, 2000), (Mooney et al., 1998).

Hearst (1992) and Sack (1994) propose cognitively inspired models for sentiment analysis. Hearst's model is inspired by Talmy's theory of force dynamics (Talmy, 1985) which describes the lexical and grammatical expressions of the interaction between two opposing entities - the Agonist and the Antagonist. Each entity expresses an intrinsic force, tending either towards motion or towards rest. The balance between these forces determines the resulting state of the interaction. In a variation on this idea, the focus is on one entity and on the events that affect it during encounters with other entities. This can be imagined as an entity following a path towards a goal or destination, and meeting barriers or facilitators along the way. This *path model* is what Hearst applies, with minor modifications, to queries that have a directional component, which imply finding whether an agent or event opposes or is in favour of another event. Sack briefly describes *SpinDoctor*, a system designed to identify the point of view of a news report. Despite having to be objective reports of facts, news reports are often biased, although sometimes not consciously. Sack's system builds on the observation that news writers are consistent in the attributes they bestow upon the actors involved in news. In order to identify the point of view of a news story, the system uses heuristics and a database of "fairy-tale-like roles" which American journalists used to describe events and participants in the first Gulf War, from (Lakoff, 1991).

Das and Chen (2001) analyse the opinions of small investors about the stock market, through messages from Yahoo!'s message board. Processing the messages downloaded from the message board involves the use of a generic English dictionary (CUOVALD), and a manually built collection of specific financial terms that manual analysis has shown to be relevant to this task. Statistical techniques help select the most discriminating words over the training data. For a specific stock, several algorithms would help suggest whether the investors' opinion is to buy or sell, or whether it is neutral.

[1]http://www.cognitrend.com

[2]http://tenorioresearch.itgo.com

Currently, sentiment analysis is approached mostly as a text classification problem. A textual unit of certain size is classified as expressing positive or negative (or favourable and unfavourable) feelings. The unit size can go from words – (Hatzivassiloglou and McKeown, 1997), (Turney and Littman, 2003), (Hatzivassiloglou and Wiebe, 2000), (Wiebe, 2000) – to full texts (of various size), starting with a small set of seed words (Turney, 2002), (Pang et al., 2002), (Pang and Lee, 2004), manually built lexicons (Subasic and Huettner, 2001), (Das and Chen, 2001), a mixture of unigrams, word sentiment measure, topic knowledge (Mullen and Collier, 2004), or even the world knowledge captured in the Open Mind Commonsense database (Liu et al., 2003).

Another way of analysing sentiment is to identify smaller text units which convey feelings and features which indicate subjective language (Wiebe, 1990), (Hatzivassiloglou and Wiebe, 2000), (Wiebe, 2000), (Wilson et al., 2004). It is known that a text, especially longer text, may express various opinions on various aspects of a product or a service. It is therefore important to separate objective from subjective text units, and proceed with sentiment analysis of the subjective parts (Pang and Lee, 2004).

Sentiment analysis is an interesting field, and research in the area grows and diversifies, as shown also in the collection (Qu et al., 2004). This special issue also contributes to the field. The four papers look at four different questions in the area of sentiment analysis: how to recognize the strength of opinions presented in texts; how important neutral examples are in classifying positive and negative examples; how valence shifters influence the sentiments expressed in a text unit; finally, what makes us laugh.

2 The papers

In accordance with the organization of the Ottawa workshop, the issue opens with a paper on recognizing the strength of opinion clauses in text. Wilson, Wiebe and Hwa continue their work on subjectivity/objectivity analysis by taking their endeavour one step further (Wilson et al., 2004). Identifying strong and weak opinion clauses will allow both people and automated systems to follow the evolution of feelings towards issues, products or services. More and more often, such opinions appear on the Web in blogs, news reports, messages posted on electronic boards and so on.

Koppel and Schler address an important, and thus far underappreciated, issue in Machine Learning related to sentiment analysis. It seems customary to focus on classifying sentiment using positive and negative examples (Pang et al., 2002),(Turney, 2002),(Turney and

Littman, 2003). There may be an implicit assumption that units which do not express feelings are irrelevant to this type of learning. Research on classification finer-grained than just positive-negative does not explicitly consider neutral examples – the approach is to compute a metric which combines label similarity with sentiment analysis to arrive at a label assignment (Pang and Lee, 2005). Koppel and Schler show that using neutral examples leads to significant improvements in learning to distinguish positive and negative examples.

Positive and negative sentiment analysis often relies on such indicators as seed words – nouns, verbs and adjectives – with a semantic orientation that people agree upon. A word's semantic orientation may diminish or even change because of so-called *valence shifters*. In the task of automatically classifying reviews as positive or negative, Inkpen and Kennedy explore three types of valence shifters: negations (they reverse the polarity of a term), intensifiers and diminishers (they affect the degree to which a term is positive or negative).

The special issue concludes on a cheerful note, with an excursion into computational humour. In more typical sentiment analysis tasks, feelings are captured in words on whose orientation people tend to agree. While a joke also relies on words to elicit a positive feeling, there is seldom one crucial funny word. We see word play, or a seemingly serious reference made amusing in context, and so on. It is not easy to identify word combinations that bring about a humorous effect. Mihalcea and Strapparava explore one-liner jokes and find out what makes them funny by comparing them with similarly worded serious sentences and equally serious proverbs.

References

Sanjiv Das and Mike Chen. 2001. Yahoo! for Amazon: Sentiment parsing from small talk on the Web. In *Proceedings of the 8th Asia Pacific Finance Association Annual Conference – APFA 2001*, Bangkok, Thailand.

Kushal Dave, Steve Lawrence, and David M. Pennock. 2003. Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In *Proceedings of the 12th international conference on World Wide Web (WWW'03)*, pages 519–528.

Vasileios Hatzivassiloglou and Kathleen McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the 35th ACL/8th EACL*, pages 174–181, Madrid, Spain.

Vasileios Hatzivassiloglou and Janyce Wiebe. 2000. Effects of adjective orientation and gradability on sentence subjectivity. In *Proceedings of COLING*, pages 299–305, Saarbrucken, Germany.

Marti Hearst. 1992. Direction-based text interpretation as an information access refinement. In Paul Jacobs, editor, *Text-based Intelligent Systems*, pages 257–274. Lawrence Erlbaum Associates.

Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the 10th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.

George Lakoff. 1991. Metaphor and war: the metaphor system used to justify the war in the gulf. *Journal of Urban and Cultural Studies*, 2(1):59–72.

Hugo Liu, Henry Lieberman, and Ted Selker. 2003. A model of textual affect sensing using real-world knowledge. In *Proceedings of the 7th International Conference on Intelligent User Interfaces*, pages 125–132.

Ray Mooney, Paul Bennett, and Loriene Roy. 1998. Book recommending using text categorization with extracted information. In *Proceedings of the AAAI workshop on Recommender Systems*, pages 70–74, Madison, WI, USA.

Tony Mullen and Nigel Collier. 2004. Sentiment analysis using support vector machines with diverse information sources. In *Proceedings of EMNLP*, pages 412–418, Barcelona, Spain.

Bo Pang and Lillian Lee. 2004. A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd ACL*, pages 271–278, Barcelona, Spain.

Bo Pang and Lillian Lee. 2005. Seeing stars: exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd ACL*, pages 115–124, Ann Arbour, MI, USA.

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment classification using machine learning techniques. In *Proceedings of EMNLP*, pages 79–86, Philadelphia, PA, USA.

Yan Qu, James G. Shanahan, and Janyce Wiebe, editors. 2004. *AAAI Spring Symposium on Exploring Attitude and Affect in Text.* AAAI Press, Stanford University.

Warren Sack. 1994. On the computation of point of view. In *Proceedings of the 12th AAAI*, pages 1488, student abstract, Seatle, WA, USA.

Pero Subasic and Alison Huettner. 2001. Affect analysis of text using fuzzy semantic typing. *IEEE-FS*, 9:483–496.

Leornard Talmy. 1985. Force dynamics in language and thought. In *Parasession on Causativity and Agentivity.* Chicago Linguistic Society (21st Regional Meeting), University of Chicago.

Loren Tarveen, Will Hill, Brian Amento, David McDonald, and Josh Creter. 1997. PHOAKS: A system for sharing recommendations. *Communications of the ACM*, 40(3):59–62.

Junichi Tatemura. 2000. Virtual reviewers for collaborative exploration of movie reviews. In *Proceedings of the 5th International Conference on Intelligent User Interfaces*, pages 272–275, New Orleans, LO, USA.

Peter Turney and Michael Littman. 2003. Measuring praise and criticism: inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21(4):315–346.

Peter Turney. 2002. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th ACL*, pages 417–424, Philadelphia, PA, USA.

Janyce Wiebe. 1990. *Recognizing Subjective Sentences: A Computational Investigation of Narrative Text*. Ph.D. thesis, State University of New York at Buffalo.

Janyce Wiebe. 2000. Learning subjective adjectives from corpora. In *Proceedings of the 7th National Conference on Artificial Intelligence and the 12th Conference on Innovative Applications of Artificial Intelligence*, pages 735–740.

Theresa Wilson, Janyce Wiebe, and Rebecca Hwa. 2004. Just how mad are you? finding strong and weak opinion clauses. In *Proceedings of the 19th National Conference on Artificial Intelligence*, pages 761–769, San Jose, CA, USA.