Analyzing Newswire and Social Media Data Using Multi-Vector Sentiment Analysis

Kelly Enochson, PhD*; Gregory Roberts; Michael Sorah; Jamie Thompson
Rosoka Software, Inc., 950 Herndon Parkway, Suite 280, Herndon, VA 20170

Abstract

The vast amount of written text available on the Internet provides a treasure trove of information for intelligence and security analysts, but only if the useful data can be quickly identified among all the irrelevant information. Many analytical tools available provide information about the sentiment of written texts; however, these tools typically utilize only one measure of sentiment in their metrics. Rosoka Software leverages psycholinguistic research across multiple sentiment vectors to provide precise information about an author's language and pinpoint documents that do not follow predicted patterns. Rosoka's multi-vector sentiment analysis uses four metrics to help analysts identify outliers in their data, recognize heightened emotional language, sort data by media type, and subset large data sets into only the documents that require further assessment.

1 Introduction

Rosoka Software's multilingual entity and relationship extraction software provides entity-level and document-level sentiment analysis across multiple vectors. Rosoka leverages unique algorithms that incorporate ***polarity***, ***mood***, ***aspect***, and ***intensity*** to provide a detailed representation of sentiment in a text. While it is not possible to use a text to determine an author's actual intent with certainty, it is possible to measure readers' responses to specific lexical items and linguistic structures within the text. Rosoka's ***mood***, ***intensity***, and ***aspect*** metrics are based on such psycholinguistic research (Bradley and Lang 1999; Medler, Arnoldussen, Binder, and Seidenberg 2005; Aquino and Arnell 2007; Janschewitz 2008; Eilola and Havelka 2010; Bestgen and Vincze 2012; Juhasz and Yap 2013; Warriner, Kuperman, and Brysbaert 2013). Leveraging these metrics together allows analysts to make informed inferences about an author's intent or purpose.

2 Sentiment Metrics

Rosoka's sentiment metrics are based on published studies that measure people's responses to, and perceptions of, various lexical items (Bradley and Lang 1999; Medler, Arnoldussen, Binder, and Seidenberg 2005; Aquino and Arnell 2007; Janschewitz 2008; Eilola and Havelka 2010; Bestgen and Vincze 2012; Juhasz and Yap 2013; Warriner, Kuperman, and Brysbaert 2013). Tens of thousands of lexical items in Rosoka's dictionaries are tagged not only with pragmatic information but also with these sentiment measures.

2.1 Polarity

Polarity is a measurement of how positive, negative, or neutral the language is about a particular entity or in the document as a whole, taking into consideration the salience of various words and phrases. This metric is calculated at both the entity and document level, and is measured on a floating point scale from −3 (negative) to +3 (positive), where 0 represents neutral polarity. For example, words like *elegant* and *strongest* have positive polarity, while

* Corresponding author: Kelly Enochson, kenochson@rosoka.com

words like *arson* and *shitty*[1] have negative polarity.

2.2 Mood

Mood is a measurement of the level of emotion of the language used to describe a particular entity in a document. This measure roughly indicates the degree of happiness/sadness associated with the language. While ***polarity*** measures the valence of the events themselves, ***mood*** measures an emotional response typically elicited in readers or listeners upon hearing the words. This metric is also computed at both the entity and document level, and represented as a floating point value on a scale from −3 (negative mood, e.g., "sad") to +3 (positive mood, e.g., "happy"). For example, words like *funny* and *affection* have a positive mood, while words like *cancer* and *racist* have a negative mood.

Polarity and ***mood*** differ in that the former is a somewhat objective measure of the information presented in the text, while the latter is a measure of the level of emotion in the language used. Taken together, these two measures could denote bias if valences are opposite. For example, if ***polarity*** is negative but ***mood*** is positive, this may indicate that the text is conveying a positive response to a negative event. An example of this is a news report about the extrajudicial killing of an individual that has a negative polarity but a positive mood. This suggests the author views the victim as a "bad guy" and is happy or relieved about his death.

2.3 Aspect

Aspect measures how controlled or in control the language in a text makes a reader feel. For example, words like *abandonment* and *abduction* have a very negative aspect because these words tend to make readers feel a loss of control. By contrast, words like *accomplish* and *motivate* have a very positive aspect because readers tend to feel in control when reading these words. ***Aspect*** is measured on a floating point scale ranging from −3 (controlled) to +3 (in control). Very positive or negative aspect values generally indicate that the author is part of the narrative, while neutral aspect values generally indicate that the author is not part of the narrative. Very negative values for aspect tend to indicate that the author is attempting to bully or dominate the reader, while very positive values tend to indicate that the author is attempting to persuade the reader.

2.4 Intensity

Intensity is a measurement of the level of activation used to describe a particular entity; in other words, whether an entity is activated/aroused or deactivated/calm. It is measured on a floating point scale ranging from 0 (no activation) to +3 (high activation). Words like *catastrophic* or *prestigious* evoke a high level of activation with a value of +3, while words like *thermometer* and *syllabus* evoke a low level of activation with a value of 0. Unlike ***polarity*** and ***mood***, ***intensity*** does not have a negative range.

Consider a use case where a company is interested in analyzing corporate emails to identify potential corporate espionage. Rosoka can analyze, for example, a large set of emails by many different authors. Using ***polarity*** as the metric of interest, individual emails by a particular author that have a polarity value outside two standard deviations from the mean can be identified as outliers. However, because outliers with a low intensity measure are typically of less concern than outliers with high intensity, the data can be further subset using this metric, allowing analysts to work more efficiently.

2.5 Sentiment Measure (S-M)

The three vectors of ***polarity***, ***mood***, and ***intensity*** are computed and transformed at the document level, providing an overall document ***sentiment measure***, or ***S-M***. Since ***aspect***

[1] In addition to understanding standard English and more formal registers, Rosoka also understands informal registers and nonstandard varieties of English (and other languages), as is crucial for sentiment analysis of social media data.

provides a directionality of action, a projection of the remaining vectors into one new vector provides a two-dimensional vector space where different regions map to a directionality and likelihood of action. The ***S-M*** is formed from the cross product of the polarity and mood vectors, using intensity as a (non-linear but locally linear) scalar on the resultant vector. Note that this is a linear transform from the original vector space to a two-dimensional space. This document-level computation of Rosoka's sentiment vectors can provide information about the data source and allow analysts to draw inferences about the attitude of the author and the likelihood that action will result from the piece of writing.

A negative S-M indicates that polarity and mood have opposite signs, i.e., one is a positive value and one is a negative value. This indicates that the author's mood does not correspond to the valence of the event. This is the case described above in the news article about an extrajudicial killing. A S-M with a high numeric value, i.e., closer to +/−3, indicates a high intensity value, while a S-M with a value closer to 0 indicates a low intensity value. Aspect can be viewed as a directional vector, with negative aspect indicating dominance and positive aspect indicating persuasion.

3 Method

The current experiment tests Rosoka's four sentiment vectors and S-M by comparing two different document sets from different sources. The first set is 500 general news documents spanning a range of topics including terrorism, cyber security, law enforcement, finance, and healthcare. The second set is 500 tweets using the hashtag #tacotruckoneverycorner. In an attempt to collect both positive and negative tweets, efforts were made to include tweets speaking positively and negatively about both Donald Trump and Hillary Clinton. These two data sets were processed using Rosoka Extraction to calculate ***polarity***, ***mood***, ***intensity***, ***aspect***, and ***S-M***. A significant difference in scores across these metrics would indicate that Rosoka's multi-vector sentiment analysis differentiates between different types of media and identify what the specific differences are.

4 Results

Independent Welch's t-tests were performed on all sentiment vectors as well as the overall S-M score. Results indicate that there is a significant difference in the polarity scores for news ($M = -0.76$, $sd = 1.19$) and Twitter ($M = -0.06$, $sd = 0.87$) data ($t(895) = -3.37$, $p = 0.001$), indicating that overall, news data tends to discuss more negative events than Twitter data. There is a marginally significant difference in mood scores for news ($M = 0.96$, $sd = 0.34$) and Twitter ($M = 1.17$, $sd = 0.71$) data ($t(700) = -1.89$, $p = 0.06$), indicating that tweets tend to have more negative emotional language than news articles. There is not a significant difference in intensity scores for news ($M = 1.51$, $sd = 0.18$) and Twitter ($M = 1.56$, $sd = 0.39$) data ($t(678) = -0.82$, $p = 0.41$), indicating that there is no difference in the level of activation of the language used in the two data sources. There is a significant difference in aspect scores between news ($M = 0.70$, $sd = 0.50$) and Twitter ($M = 1.10$, $sd = 0.88$) data ($t(775) = -2.81$, $p = 0.006$), indicating that Twitter data is more likely to use language attempting to persuade the reader. Median scores and distributions for the four sentiment metrics are shown in **Figure 1**.

Finally, there is a significant difference in S-M scores between news ($M = -0.57$, $sd = 0.96$) and Twitter ($M = 0.53$, $sd = 0.94$) data ($t(979) = -5.76$, $p > 0.0001$). The mean news S-M score is negative, which, as discussed above, indicates that the polarity and mood measures, on average, have opposite signs. This is as expected, since news data tends to report on negative events without emotional language. On the other hand, the mean Twitter S-M score is positive, indicating that, in general, the emotional language in tweets tends to match the polarity of the events discussed.

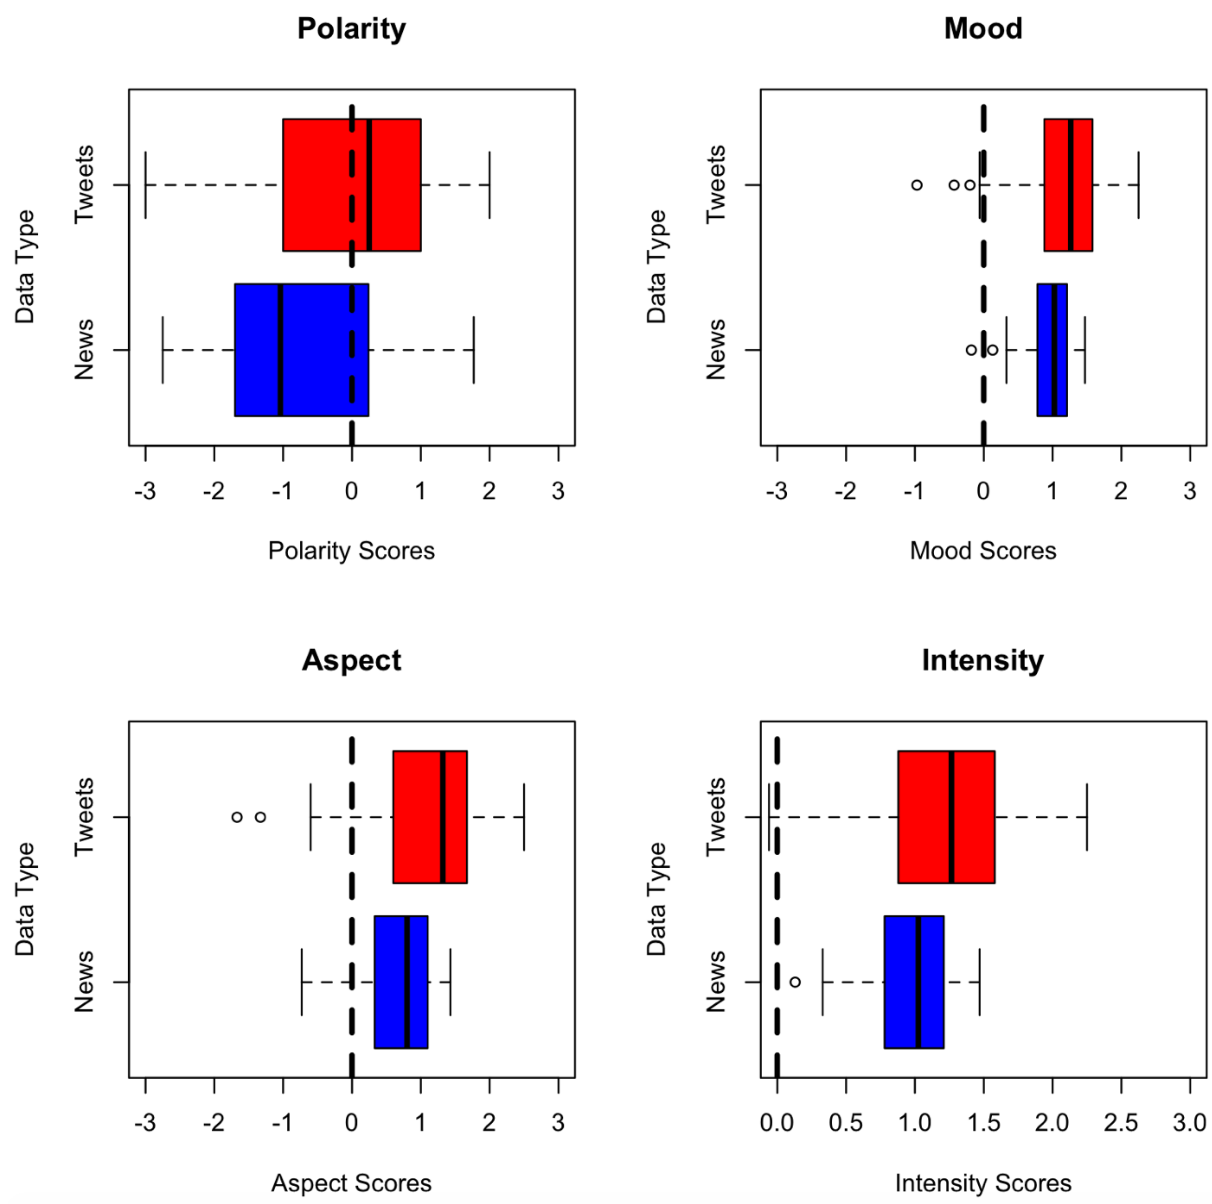

Figure 1: Boxplots of the four sentiment analysis metrics. Dotted lines indicate neutral values; solid lines indicate median scores; colored areas indicate interquartile ranges; whiskers indicate ranges; open circles indicate outliers.

The difference in S-M between these two data sources is primarily driven by differences in polarity and aspect, with news data being more negative and less persuasive, as one might expect. For all vectors except polarity, Twitter data has a larger standard deviation than news data, indicating greater variation among tweets than among news articles. This information could be used to identify specific tweets of interest, or news articles that look like tweets based on sentiment metrics. These documents in particular would likely be of greater interest to analysts.

5 Conclusion

This experiment shows that Rosoka's multi-vector sentiment analysis is able to differentiate between different media sources, in this case news and Twitter. The difference in S-M between these two data sources is primarily driven by differences in polarity and aspect, with news data being more negative and less persuasive. These metrics can help analysts identify outliers in their data, recognize heightened emotional language, sort data by type, and subset large data sets into only the documents that require further assessment. Further research is necessary to determine the true impact of documents that lie outside the expected vector projections of their document type, but we hypothesize that those data could prove meaningful with a better understanding the true pragmatic nature of the author's intended message.

References

Agarwal, A., Xie, B., Vovsha, I., Rambow, O., & Passonneau, R. (2011). Sentiment analysis of twitter data. In *Proceedings of the workshop on languages in social media* (pp. 30-38). Association for Computational Linguistics.

Aquino, J. M., & Arnell, K. M. (2007). Attention and the processing of emotional words: Dissociating effects of arousal. *Psychonomic Bulletin & Review*, *14*(3), 430-435.

Bestgen, Y., & Vincze, N. (2012). Checking and bootstrapping lexical norms by means of word similarity indexes. *Behavior research methods*, *44*(4), 998-1006.

Bhatia, P., Ji, Y., & Eisenstein, J. (2015). Better document-level sentiment analysis from rst discourse parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, (EMNLP)*.

Boiy, E., Hens, P., Deschacht, K., & Moens, M. F. (2007). Automatic Sentiment Analysis in On-line Text. In *ELPUB* (pp. 349-360).

Bosco, C., Patti, V., & Bolioli, A. (2015). Developing Corpora for Sentiment Analysis: The Case of Irony and Senti-TUT. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): Instruction manual and affective ratings* (pp. 1-45). Technical report C-1, the center for research in psychophysiology, University of Florida.

Cambria, E., Schuller, B., Xia, Y., & Havasi, C. (2013). New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems*, (2), 15-21.

Choi, Y., & Cardie, C. (2008). Learning with compositional semantics as structural inference for subsentential sentiment analysis. In*Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 793-801). Association for Computational Linguistics.

Eilola, T. M., & Havelka, J. (2010). Affective norms for 210 British English and Finnish nouns. Behavior Research Methods, 42(1), 134-140.

Godbole, N., Srinivasaiah, M., & Skiena, S. (2007). Large-Scale Sentiment Analysis for News and Blogs. *ICWSM*, *7*(21), 219-222.

Hu, X., Tang, J., Gao, H., & Liu, H. (2013). Unsupervised sentiment analysis with emotional signals. In *Proceedings of the 22nd international conference on World Wide Web* (pp. 607-618). International World Wide Web Conferences Steering Committee.

Hu, X., Tang, L., Tang, J., & Liu, H. (2013). Exploiting social relations for sentiment analysis in microblogging. In *Proceedings of the sixth ACM international conference on Web search and data mining* (pp. 537-546). ACM.

Hutto, C. J., & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International AAAI Conference on Weblogs and Social Media*.

Janschewitz, K. (2008). Taboo, emotionally valenced, and emotionally neutral word norms. *Behavior Research Methods*, *40*(4), 1065-1074.

Juhasz, B. J., & Yap, M. J. (2013). Sensory experience ratings for over 5,000 mono-and disyllabic words. *Behavior Research Methods*, *45*(1), 160-168.

Kiritchenko, S., Zhu, X., & Mohammad, S. M. (2014). Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*, 723-762.

Kumar, A., & Sebastian, T. M. (2012). Sentiment analysis on twitter. *IJCSI International Journal of Computer Science Issues*, *9*(4), 372-373.

Kumar, A., & Sebastian, T. M. (2012). Sentiment analysis on twitter. *IJCSI International Journal of Computer Science Issues*, *9*(4), 372-373.

Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, *5*(1), 1-167.

Maynard, D., & Greenwood, M. A. (2014). Who cares about Sarcastic Tweets? Investigating the Impact of Sarcasm on Sentiment Analysis. In *LREC* (pp. 4238-4243).

McDonald, R., Hannan, K., Neylon, T., Wells, M., & Reynar, J. (2007). Structured models for fine-to-coarse sentiment analysis. In *Annual Meeting-Association For Computational Linguistics* (Vol. 45, No. 1, p. 432).

Medler, D. A., Arnoldussen, A., Binder, J.R., & Seidenberg, M.S. (2005). *The Wisconsin Perceptual Attribute Ratings Database*. http://www.neuro.mcw.edu/ratings/

Morency, L. P., Mihalcea, R., & Doshi, P. (2011, November). Towards multimodal sentiment analysis: Harvesting opinions from the web. In *Proceedings of the 13th international conference on multimodal interfaces*(pp. 169-176). ACM.

Mukherjee, S., & Bhattacharyya, P. (2012). Feature specific sentiment analysis for product reviews. In *Computational Linguistics and Intelligent Text Processing* (pp. 475-487). Springer Berlin Heidelberg.

Nasukawa, T., & Yi, J. (2003, October). Sentiment analysis: Capturing favorability using natural language processing. In *Proceedings of the 2nd international conference on Knowledge capture* (pp. 70-77). ACM.

Ortigosa-Hernández, J., Rodríguez, J. D., Alzate, L., Lucania, M., Inza, I., & Lozano, J. A. (2012). Approaching Sentiment Analysis by using semi-supervised learning of multi-dimensional classifiers. *Neurocomputing*, *92*, 98-115.

Osimo, D., & Mureddu, F. (2012). Research challenge on opinion mining and sentiment analysis. *Universite de Paris-Sud, Laboratoire LIMSI-CNRS, Bâtiment*, *508*.

Pak, A., & Paroubek, P. (2010, May). Twitter as a Corpus for Sentiment Analysis and Opinion Mining. In *LREc* (Vol. 10, pp. 1320-1326).

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, *2*(1-2), 1-135.

Prabowo, R., & Thelwall, M. (2009). Sentiment analysis: A combined approach. *Journal of Informetrics*, *3*(2), 143-157.

Rambocas, M., & Gama, J. (2013). *Marketing research: The role of sentiment analysis* (No. 489). Universidade do Porto, Faculdade de Economia do Porto.

Saif, H., Fernandez, M., He, Y., & Alani, H. (2013). Evaluation datasets for Twitter sentiment analysis: a survey and a new dataset, the STS-Gold.

Saif, H., Fernández, M., He, Y., & Alani, H. (2014). On stopwords, filtering and data sparsity for sentiment analysis of Twitter.

Saif, H., Fernandez, M., He, Y., & Alani, H. (2014). Senticircles for contextual and conceptual semantic sentiment analysis of twitter. In *The Semantic Web: Trends and Challenges* (pp. 83-98). Springer International Publishing.

Shukla, A. (2011). Sentiment Analysis of Document Based on Annotation.*arXiv preprint arXiv:1111.1648*.

Sindhwani, V., & Melville, P. (2008). Document-word co-regularization for semi-supervised sentiment analysis. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on* (pp. 1025-1030). IEEE.

Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational linguistics*, *37*(2), 267-307.

Vinodhini, G., & Chandrasekaran, R. M. (2012). Sentiment analysis and opinion mining: a survey. *International Journal*, *2*(6).

Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior research methods*, *45*(4), 1191-1207.

Westerski, A. (2007). Sentiment Analysis: Introduction and the State of the Art overview. *Universidad Politecnica de Madrid, Spain*, 211-218.

Wilson, T., Wiebe, J., & Hoffmann, P. (2005). Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on human language technology and empirical methods in natural language processing* (pp. 347-354). Association for Computational Linguistics.

Yang, B., & Cardie, C. (2014). Context-aware Learning for Sentence-level Sentiment Analysis with Posterior Regularization. In *ACL (1)* (pp. 325- 335).

Yi, J., Nasukawa, T., Bunescu, R., & Niblack, W. (2003). Sentiment analyzer: Extracting sentiments about a given topic using natural language processing techniques. In *Data Mining, 2003. ICDM 2003. Third IEEE International Conference on* (pp. 427-434). IEEE.

Zhao, J., Dong, L., Wu, J., & Xu, K. (2012). Moodlens: an emoticon-based sentiment analysis system for Chinese tweets. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 1528-1531). ACM.