Identifying Opinionated Sentences

Theresa Wilson
Intelligent Systems Program
University of Pittsburgh
twilson@cs.pitt.edu

David R. Pierce
Department of Computer Science
and Engineering
University of Buffalo
The State University of New York
drpierce@cse.buffalo.edu

Janyce Wiebe
Department of Computer Science
University of Pittsburgh
wiebe@cs.pitt.edu

1 Introduction

Natural language processing applications that summarize or answer questions about news and other discourse need to process information about opinions, emotions, and evaluations. For example, a question answering system that could identify opinions in the news could answer questions such as the following:

> Was the 2002 presidential election in Zimbabwe regarded as fair?

> What was the world-wide reaction to the 2001 annual U.S. report on human rights?

In the news, editorials, reviews, and letters to the editor are sources for finding opinions, but even in news reports, segments presenting objective facts are often mixed with segments presenting opinions and verbal reactions. This is especially true for articles that report on controversial or "lightning rod" topics. Thus, there is a need to be able to identify which sentences in a text actually contain expressions of opinions and emotions.

We demonstrate a system that identifies opinionated sentences. In general, an opinionated sentence is a sentence that contains a significant expression of an opinion, belief, emotion, evaluation, speculation, or sentiment. The system was built using data and other resources from a summer workshop on multi-perspective question answering (Wiebe et al., 2003) funded under ARDA NRRC.[1]

2 Opinion Recognition System

2.1 System Architecture

The opinion recognition system takes as input a URL or raw text document and produces as output an HTML version of the document with the opinionated sentences found by the system highlighted in bold. Figure 2.1 shows a news article that was processed by the system.

When the opinion recognition system receives a document, it first uses GATE (Cunningham et al., 2002) (modified to run in batch mode) to tokenize, sentence split, and part-of-speech tag the document. Then the document is stemmed and searched for features of opinionated language. Finally, opinionated sentences are identified using the features found, and they are highlighted in the output.

2.2 Features

The system uses a combination of manually and automatically identified features. The manually identified features were culled from a variety of sources, including (Levin, 1993) and (Framenet, 2002). In addition to features learned in previous work (Wiebe et al., 1999; Wiebe et al., 2001), the automatically identified features include new features that were learned using information extraction techniques (Riloff and Jones, 1999; Thelen and Riloff, 2002) applied to an unannotated corpus of world news documents.

2.3 Evaluation

We evaluated the system component that identifies opinionated sentences on a corpus of 109 documents (2200 sentences) from the world news. These articles were annotated for expressions of opinions as part of the summer workshop on multi-perspective question answering. In this test corpus, 59% of sentences are opinionated sentences. By varying system settings, the opinionated sentence recognizer may be tuned to be very precise (91% precision), identifying only those sentences it is very sure
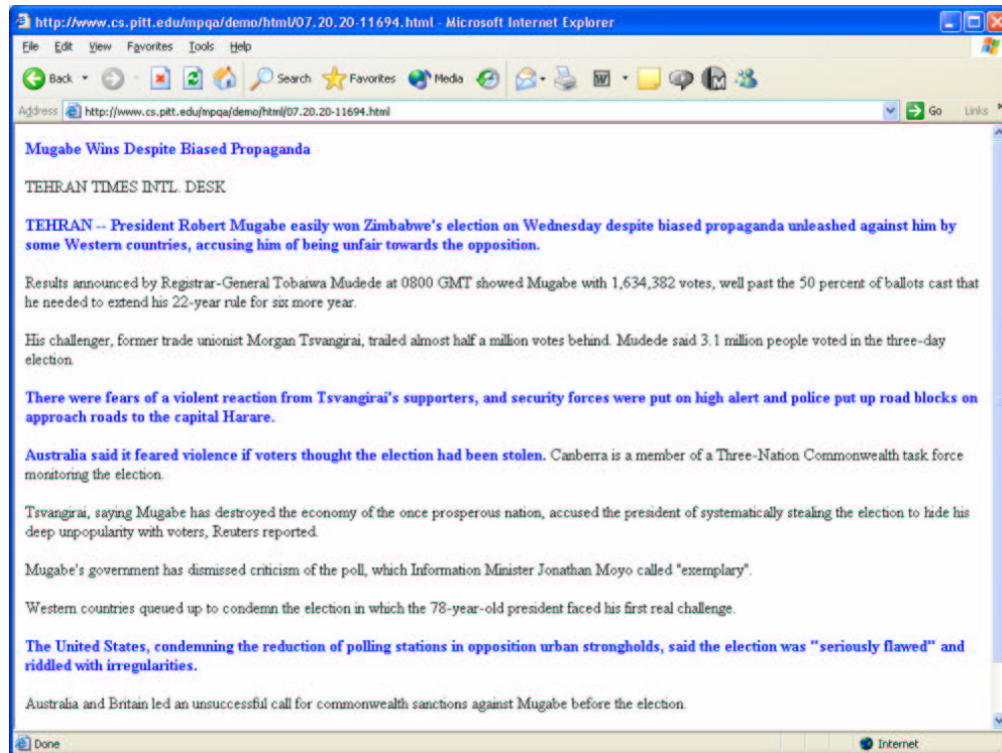
Figure 1: Example of an article processed by the opinionated sentence recognition system. Sentences identified by the system are highlighted in bold.

are opinionated (33% recall), or less precise (82% precision), identifing many more opinionated sentences (77% recall), but also making more errors.

References

Hamish Cunningham, Diana Maynard, Kalina Bontcheva, and Valentin Tablan. 2002. GATE: A framework and graphical development environment for robust nlp tools and applications. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*.

Framenet. 2002. *http://www.icsi.berkeley.edu/~framenet/*.

Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press, Chicago.

E. Riloff and R. Jones. 1999. Learning Dictionaries for Information Extraction by Multi-Level Bootstrapping. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*.

M. Thelen and E. Riloff. 2002. A bootstrapping method for learning semantic lexicons using extraction pattern contexts. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*.

J. Wiebe, R. Bruce, and T. O'Hara. 1999. Development and use of a gold standard data set for subjectivity classifications. In *Proc. 37th Annual Meeting of the Assoc. for Computational Linguistics (ACL-99)*, pages 246–253, University of Maryland, June. ACL.

J. Wiebe, T. Wilson, and M. Bell. 2001. Identifying collocations for recognizing opinions. In *Proc. ACL-01 Workshop on Collocation: Computational Extraction, Analysis, and Exploitation*, July.

J. Wiebe, E. Breck, C. Buckley, C. Cardie, P. Davis, B. Fraser, D. Litman, D. Pierce, E. Riloff, T. Wilson, D. Day, and M. Maybury. 2003. Recognizing and organizing opinions expressed in the world press. In *Working Notes - New Directions in Question Answering (AAAI Spring Symposium Series)*.