Verb-Triggered Event Detection and Classification (Extended Abstract)

Eric Bell, Liam McGrath, and **Michelle Gregory**
Pacific Northwest National Laboratory

1 Problem Description

In analyzing natural language text, it is of interest to be able to recognize and correctly classify events that occur. An event is a specific occurrence of some happening that involves one or more participants. Event analysis typically involves three tasks: determining the span of text identifying an event, determining the event trigger, and determining the participants in an event. In this paper, we concentrate on detecting events triggered by words and then correctly classifying those events by type.

2 Related Work

Current techniques for event detection and classification include both rule-based and statistical learning approaches. In either type of approach, similar features are used for event detection. Character affixes, nominalization suffix, part of speech, light verb, subject syntactic category, morphological stem, verb root, and WordNet hypernyms have all been used to classify events (Jurafsky and Martin, 2008). The TimeBank corpus is frequently used to explain and show annotations of events (Pustejovsky et al., 2003).

Traditional event detection used a co-occurrence technique based on the similarity of features between seen documents and a new unseen document (Wei and Lee, 2004). This technique led to problems with orientation and word mismatch. As a result, traditional techniques are limited in application (Wei and Lee, 2004). In more recent years, studies have concentrated on improving feature-based event detection by using a combination of text categorization and information extraction techniques (Wei and Lee, 2004). As an example, event classification is often based on and extended to a classification of the entire sentence (Naughton et al., 2008) (Chieu and Lee, 2004).

Event detection is often linked with temporal organization; therefore, many modern approaches leverage temporal aspects to achieve better results. Both supervised-learning and unsupervised-clustering algorithms for document classification then use that agglomerative clustering for retrospective event detection (Yang et al., 1999). While such approaches are effective and useful, the domain is limited to document sets that have temporal aspects that are useful in the classification of events within a document. Typical text does not fall within that domain.

The most recent work on event extraction used a nearest neighbor classifier to distinguish positive and negative instances of events, and a subsequent nearest neighbor classifier to determine event class (Ahn, 2006).

3 Methodology

While many techniques are proposed for identifying taggable events in text, we explore the use of sense disambiguation techniques for event classification. The problem can be thought of in the following manner: given a trigger or word, there are two possibilities; the trigger (verb) can represent an event of interest, or it can represent an event or something else in which we have no interest. A single word can trigger one or more event types. For example, "transfer" could refer to transferring ownership of an item, transferring money, or transferring personnel from one location to another. Each sense of the word is linked with an event type. Therefore, a trigger can trigger one or more event classes that can be thought of as "senses" of the word. One "sense" corresponds to the word being not of interest, and each of the other "senses" correspond to event classes of interest. Then, the event classification becomes a matter of sense disambiguation of the word. For example, to determine the event class of an occurrence of the word *transfer*, determine the "sense" of *transfer* being used (Person, Money, Ownership, etc.). Our work is distinct from previous attempts at this problem in the use of single classifier to handle both identification and classification in a single step. This is achieved by considering negative instances of event triggers as a class of their own.

4 Implementation

To collect baseline results, we started with a rule-based system that mapped verbs to the most relevant event class. From that baseline, we trained Lingpipe's TF/IDF model to recognize and classify instances of events. This consisted of developing a vocabulary from all tagged verbs in the training corpus and then training on feature vectors involving the surrounding context for all occurrences of those verbs with the appropriate classification. We then used the trained classifier on every verb in the testing data, receiving the resultant classifications. The use of the classifier showed significant improvement over the previously used rule-based system as evidenced by an increase in F-Measure from 0.31 to 0.69.

5 Future Work

In the future, we intend to extend our approach to handle events triggered by other parts of speech, not just verbs. We would like to look at the difference between using a classifier for each part of speech with a single classifier used for all parts of speech. We intend to use the automatic content extraction contest data and scoring method to compare our system with published system results. Lastly, we will compare our system with the previously implemented system of separate event identification and event classification steps.

References

David Ahn. 2006. The stages of event extraction. In *ARTE '06: Proceedings of the Workshop on Annotating and Reasoning about Time and Events*, pages 1–8, Morristown, NJ, USA. Association for Computational Linguistics.

Hai Leong Chieu and Yoong Keok Lee. 2004. Query based event extraction along a timeline. In *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 425–432, New York, NY, USA. ACM.

Daniel Jurafsky and James H. Martin. 2008. *Speech and Language Processing (2nd Edition) (Prentice Hall Series in Artificial Intelligence)*. Prentice Hall, 2 edition, May.

Martina Naughton, Nicola Stokes, and Joe Carthy. 2008. Investigating statistical techniques for sentence-level event classification. In *COLING '08: Proceedings of the 22nd International Conference on Computational Linguistics*, pages 617–624, Morristown, NJ, USA. Association for Computational Linguistics.

J. Pustejovsky, P. Hanks, R. Sauri, A. See, R. Gaizauskas, A. Setzer, D. Radev, B. Sundheim, D. S. Day, L. Ferro, and M. Lazo. 2003. The timebank corpus. In *Proceedings of Corpus Linguistics 2003 Conference*, pages 647–656. UCREL.

Chih-Ping Wei and Yen-Hsien Lee. 2004. Event detection from online news documents for supporting environmental scanning. *Decis. Support Syst.*, 36(4):385–401.

Yiming Yang, Jaime G. Carbonell, Ralf D. Brown, Thomas Pierce, Brian T. Archibald, and Xin Liu. 1999. Learning approaches for detecting and tracking news events. *IEEE Intelligent Systems*, 14(4):32–43.