

## Using computers to analyse sentiments

# An emotional response

## Software that can tell when people are getting upset

Oct 6th 2009 | From *The Economist* online

THE difference between saying what you mean and meaning what you say is obvious to most people. To computers, however, it is trickier. Yet getting them to assess intelligently what people mean from what they say would be useful to companies seeking to identify unhappy customers and intelligence agencies seeking to identify dangerous individuals from comments they post online.

Computers are often inept at understanding the meaning of a word because that meaning depends on the context in which the word is used. For example "killing" is bad and "bacteria" are bad but "killing bacteria" is often good (unless, that is, someone is talking about the healthy bacteria present in live yogurt, in which case, it would be bad).

An attempt to enable computers to assess the emotional meaning of text is being led by Stephen Pulman of the University of Oxford and Karo Moilanen, one of his doctoral students. It uses so-called "sentiment analysis" software to assess text. The pair have developed a classification system that analyses the grammatical structure of a piece of text and assigns emotional labels to the words it contains, by looking them up in a 57,000-word "sentiment lexicon" compiled by people. These labels can be positive, negative or neutral. Words such as "never", "failed" and "prevent" are tagged as "changing" or "reversive" words because they reverse the sentiment of word they precede.

The analysis is then broken into steps that progressively take into account larger and larger grammatical chunks, updating the sentiment score of each entity as it goes. The grammatical rules determine the effect of one chunk of text on another. The simplest rule is that positive and negative sentiments both overwhelm neutral ones. More complex syntactic rules govern seemingly conflicting cases such as "holiday hell" or "abuse helpline" that make sense to people but can confuse computers.

By applying and analysing emotional labels, the software can construct sentiment scores for the concepts mentioned in the text, as a combination of positive, negative and neutral results. For example, in the sentence, "The region's largest economies were still mired in recession," the parser finds four of the words in the sentiment lexicon: largest (positive, neutral or negative); economies (positive or neutral); mired (negative); and recession (negative). It then analyses the sentence structure, starting with "economies" and progressing to "largest economies", "region's largest economies" and "the region's largest economies". At each stage, it computes the changing sentiment of the sentence. It then does the same for the second half of the sentence.

Instead of simply adding up the number of positive and negative mentions for each concept, the software applies a weighting to each one. For example, short pieces of text such as "region" are given less weight than longer ones such as "the region's largest economies". Once the parser has reassembled the original text ("the region's largest economies were still mired in recession") it can correctly identify the sentence as having a mainly negative meaning with respect to the concept of "economies".

The researchers say this approach is better than existing text-mining systems, which can inform companies what people think of them but cannot make complex links between statements. It may be, for example, that a carmaker's customers love its latest sports car, but loathe its sound system.

As well as companies seeking to better understand their customer, intelligence agencies are also becoming interested in the sentiment analysis. Some agencies are using tools developed at **the University of Arizona's artificial intelligence laboratory** to map intense, violent emotion in online forums frequented by political radicals in order to identify surges of bad feelings and even potential terrorists. Hsinchun Chen, director of the laboratory, says that out of 8m postings, his software might be able to isolate hundreds of postings by just 20 or 30 individuals that warrant a closer look. But the software can only supplement human judgment—not least because people don't always mean what they say.

