

“Treating paraphrases in the context of a training simulation”

Mary Gardiner

June 13, 2008

When either a human author or a computer natural language generation system tries to express an idea, there is usually more than one way to say it. There may be multiple words that mean very similar things, or multiple phrases that express the same idea. However, it is argued that there are no true synonyms, that even words that have very similar meanings cannot be substituted for each other in all circumstances. Automatic natural language generation systems therefore have a use for modules which make effective word and phrase choices among closely related alternatives.

In this thesis I consider the specific problem of choosing an appropriate word or phrase where the alternatives are closely related in meaning but differ in sentiment or attitude. One example is *stingy* and *frugal*, one of which is critical of what it describes and the other of which is complimentary.

The thesis will address three aspects of the problem. The first question is whether existing methods to predict word choice among closely related words are sufficient for choosing between words that differ in sentiment. There are two methods in the literature for this, both relying on statistical models of words in context. The relatively poor performance of these methods has been used to argue that statistical methods are not suitable for this task. Using a more comprehensive set of data for this thesis, I show that sets of words that differ in sentiment behave in a distinct fashion, suggesting that they are amenable to statistical approaches.

The second aspect of my research into choosing between related words or phrases that differ in sentiment is investigating whether or not including some global information about the entire text is useful in predicting word choice. I hypothesise that information about the sentiment of a document as a whole (for example, if the document is a movie review, whether it is favourable or not) will assist in choosing between closely related words that differ in sentiment. I present several models incorporating such information: a baseline model, an oracle model showing the maximum performance possible, and three models exploring actual performance when sentiment cues need to be inferred from the text.

The third aspect is an investigation into using word and phrase choice in a particular generation task: that of changing an existing text so that it is written in a different style. This extends the earlier work from words to entire paraphrases, acquiring and using paraphrases that change the style and sentiment of a text without its meaning.