

What are Subjectivity, Sentiment, and Affect?

Eduard Hovy

USC Information Sciences Institute (ISI)

Marina del Rey, CA, USA

<http://www.isi.edu/~hovy>

hovy@isi.edu

Abstract

Pragmatics—the aspects of text that signal interpersonal and situational information, complementing semantics—has been almost totally ignored in Natural Language Processing. But in the past five to eight years there has been a surge of research on the general topic of ‘opinion’, also called ‘sentiment’. Generally, research focuses on the determining the author’s opinion/sentiment about some topic within a given fragment of text. Since opinions may differ, it is granted that the author’s opinion is ‘subjective’, and the effectiveness of an opinion-determination system is measured by comparing against a gold-standard set of human annotations.

But what does ‘subjectivity’ actually mean? What are ‘opinion’ and ‘sentiment’? Lately, researchers are also starting to talk about ‘affect’, and even ‘emotion’. What are these notions, and how do they differ from one another?

Unfortunately, a survey of the research done to date shows a disturbing lack of clarity on these questions. Very few papers bother to define their terms, but simply take a set of valences such as Good–Neutral–Bad to be sufficient. More recent work acknowledges the need to specify what the opinion actually applies to, and attempts also to determine the theme. Lately, several identify the holder of the opinion. Some even try to estimate the strength of the expressed opinion.

The trouble is, the same aspect of the same object can be considered Good by one person and Bad by another, and we can often understand both their points of view. There is much more to opinion/sentiment than simply matching words and phrases that attach to the theme, and computing a polarity score. People give reasons why they like or dislike something, and these reasons pertain to their goals and plans in the case of

opinions) or their deeper emotional states (in the case of affect).

In this talk I outline a model of sentiment/opinion and of affect, and show that they appear in text in a fairly structured way, with various components. I show how proper understanding requires the reader to build some kind of person profile of the author, and claim that for systems to do adequate understanding of sentiments, opinions, and affects, they will need to do so as well. This is not a trivial challenge, and it opens the door to a whole new line of research with many fascinating and practical aspects.

About The Speaker

Dr. Hovy currently holds several positions:

- ❖ Director of the [Natural Language Group](#) at Information Sciences Institute (ISI) of the University of Southern California.
- ❖ Deputy Director of the [Intelligent Systems Division of ISI](#).
- ❖ Research Associate Professor of [Computer Science at USC](#).
- ❖ Director of the [Center for Knowledge Integration and Discovery \(CKID\)](#).
- ❖ Director of Research for the [Digital Government Research Center \(DGRC\)](#).
- ❖ Regular High-Level Visiting Scientist, International Guest Academic Talents (IGAT) Program for the Development of University Disciplines in China (111 Program), Jan 2008–Dec 2012.
- ❖ Advisory Professor at the [Beijing University of Posts and Telecommunications](#), Beijing, China.
- ❖ Concurrent Professor at the [University of Shenyang, China](#), Oct 2008–Sep 2011.