

Opinion Mining: Exploiting the Sentiment of the Crowd

Monday 21 May 2012, 9am - 1pm

Presenters:

Diana Maynard

Tutorial Programme/Overview

This tutorial will introduce the concepts of opinion mining and sentiment analysis from unstructured text, looking at why they are useful and what tools and techniques are available. We will describe a variety of general rule-based and machine learning techniques, provide some background information on the key underlying NLP processes required, and focus specifically on some of the major problems and solutions, such as detection of sarcasm, use of informal language, spam opinion detection, trustworthiness of opinion holders, and so on. We will also discuss problems associated with opinion detection in social media such as blogs, forum posts, twitter, etc. The techniques will be demonstrated with key open-source tools and applications.

Tutorial Description/Outline/Contents

The tutorial will be divided into 3 sections, as follows:

1. **Introduction to opinion mining.** This part of the tutorial will explain the motivation for such tools, introduce the main components needed for an application, and outline the major difficulties. It will look at some existing opinion mining tools and discuss their strengths and weaknesses.
2. **Applications.** The second section will explore further the concept of opinion mining and take an in-depth look at some real opinion mining applications, covering both rule-based and machine-learning approaches, in the domains of product reviews, political opinions, and discussions of social events.
3. **Challenges for social media.** Finally we will introduce and discuss some of the major challenges for working with social media, such as tweets, blogs and facebook posts, using real examples. We will also focus on some of the linguistic challenges stemming from these such as negatives, conditional statements, incorrect English, slang and swear words, irony and sarcasm, and so on.