Multimodal Sentiment Analysis
(Abstract of Invited Talk)

Rada Mihalcea

Department of Computer Science and Engineering

University of North Texas

P. O. Box 311366

Denton, TX 76203-6886, U.S.A.

rada@cs.unt.edu

Abstract

With more than 10,000 new videos posted online every day on social websites such as YouTube and Facebook, the internet is becoming an almost infinite source of information. One important challenge for the coming decade is to be able to harvest relevant information from this constant flow of multimodal data. In this talk, I will introduce the task of multimodal sentiment analysis, and present a method that integrates linguistic, audio, and visual features for the purpose of identifying sentiment in online videos. I will first describe a novel dataset consisting of videos collected from the social media website YouTube, which were annotated for sentiment polarity. I will then show, through comparative experiments, that the joint use of visual, audio, and textual features greatly improves over the use of only one modality at a time. Finally, by running evaluations on datasets in English and Spanish, I will show that the method is portable and works equally well when applied to different languages.

This is joint work with Veronica Perez-Rosas and Louis-Philippe Morency.

1