

Having the possibility of automatically tracking a person's popularity in the newspapers is an idea appealing not just to those in the media spotlight. While sentiment (subjectivity) analysis is a rapidly growing subfield of computational linguistics, no data from the news domain are yet available for Czech. We have therefore started building a manually annotated polarity corpus of sentences from Czech news texts; however, these texts have proven themselves rather unwieldy for such processing. We have also designed a classifier which should be able to track popularity based on this corpus; the classifier has been tested on a corpus of product reviews of domestic appliances and some introductory testing has been done on the nascent news corpus. As a model, we simply extract a unigram polarity lexicon from the data. We then use three related methods for identifying lemma polarity and a number of simple filters for feature selection. On the domestic appliance data, our simplest model has achieved results comparable to the state of the art, however, the properties of Czech news texts and preliminary results hint a more linguistically oriented approach might be preferable.