

Abstract: Two approaches to generating domain-specific sentiment lexicons are proposed: (i) direct generation and (ii) adaptation. The first is founded on the idea of generating a dedicated lexicon directly from labelled source data. The second approach is founded on the idea of using an existing general purpose lexicon and adapting this so that it becomes a specialised lexicon with respect to some domain. The operation of the two approaches is illustrated using a political opinion mining domain and evaluated using a large corpus of labelled political speeches extracted from political debates held within the UK Houses of Commons.

Keywords: Domain-Specific Sentiment Lexicons, Opinion Mining, Sentiment Analysis.