Abstract

Authorship Attribution is the study of identifying people by their writing style. We present several approaches to achieving this goal, looking at both Authorship Identification problems (in which we attempt to predict which of a limited set of candidate authors wrote a disputed text) and Authorship Verification problems (in which we attempt to predict whether or not two texts are written by the same author). We test and compare several techniques, including an unsupervised method that relies on descriptive statistics; methods which use Support Vector Machines; and methods which use Neural Networks. We compare our methods to previous state of the art results and present results on a new large dataset, built from Yelp reviews, which we introduce. Although in many cases our methods failed to beat previous state of the art results, we show that all three broad techniques are viable strategies for Authorship Attribution tasks, and we discuss the advantages and disadvantages of each approach.