

Confidence Estimation in Structured Prediction

Avihai Mejer

Department of Computer Science
Technion-Israel Institute of Technology
Haifa 32000, Israel
`amejer@tx.technion.ac.il`

Koby Crammer

Department of Electrical Engineering
Technion-Israel Institute of Technology
Haifa 32000, Israel
`koby@ee.technion.ac.il`

Abstract

Structured classification tasks such as sequence labeling and dependency parsing have seen much interest by the Natural Language Processing and the machine learning communities. Several online learning algorithms were adapted for structured tasks such as Perceptron, Passive-Aggressive and the recently introduced Confidence-Weighted learning. These online algorithms are easy to implement, fast to train and yield state-of-the-art performance. However, unlike probabilistic models like Hidden Markov Model and Conditional random fields, these methods generate models that output merely a prediction with no additional information regarding confidence in the correctness of the output. In this work we fill the gap proposing few alternatives to compute the confidence in the output of non-probabilistic algorithms. We show how to compute confidence estimates in the prediction such that the confidence reflects the probability that the word is labeled correctly. We then show how to use our methods to detect mislabeled words, trade recall for precision and active learning. We evaluate our methods on four noun-phrase chunking and named entity recognition sequence labeling tasks, and on dependency parsing for 14 languages.

1 Introduction

Large scale natural language processing systems are often composed of few components each designed for solving a specific task. Example tasks are part-of-speech (POS) tagging (annotate words with their grammatical role), noun-phrase (NP) chunking (identify noun-phrases), information-extraction (IE) or