Abstract

Accidental mistakes, annotator biases and encoding problems are just some of the reasons why the labels ML classifiers train on are rarely fully reliable. This thesis explores in depth a newly proposed solution to label noise, the Noisy Label Neural Network algorithm. To this algorithm, noise is a channel that transforms each correct label into either a copy or an error. The likelihood of mislabeling is given by the noise parameter, a confusion matrix which holds for each class the likelihood of being transformed into another class. The Noisy Label Neural Network simultaneously learns the noise parameter and the neural network parameters through back-propagation and Expectation Maximization algorithm iterations. Handling noise this way is successful on MNIST handwriting recognition and TIMIT phoneme classification data that has been injected with noise, but the method has not been tested on higher-level linguistic tasks. Because of the known presence of label noise in topic classification databases, I test the method on Reuters-21578 data, on which it achieves significantly better results than a standard neural network approach. Label noise is also a known problem in knowledge-based Named Entity Recognition (NER) approaches. I therefore propose a cycle of automatic gazetteer-reliant labeling, deep learning and iterative noise-distribution estimation to quickly and easily produce more reliable labels using only gazetteers. This method is very effective on uniformly noisy NER-tags and has some benefit when used on automatically generated labels. This proves noise-robust methods such as the NLNN algorithm can help deal with random errors as well as less predictable noise sources that hinder progress in Natural Language Processing.