

Student Research Symposium 2003

Satanjeev (Bano) Banerjee

Improving Language Models by Learning from Speech Recognition Errors in a Reading Tutor that Listens

The goal of this research is to use speech recognition errors to improve language model generation. This research is done within the context of project LISTEN's Reading Tutor which uses speech recognition to listen to children read stories one sentence at a time, and detects reading errors. Since the sentence the child is attempting to read is known ahead of time, the speech recognizer uses a separate language model for each sentence that consists of only the words in the sentence, and a few other words or phoneme sequences that the child is predicted to utter. We construct and evaluate an algorithm that learns to improve on such language models.

Specifically, given a set of transcribed utterances and their corresponding language models, we recognize the utterances, and compute a measure of tracking error rate by comparing the output hypotheses with the transcripts and the target text. Next we identify those n-grams in the language models that contributed to an increase or decrease in tracking error. Finally we create a feature vector for each such n-gram, and then induce a classifier that outputs the probability that a given n-gram will decrease tracking error. Given this classifier, we then implement a simple "nudge" algorithm that increases or decreases the transition probabilities of n-grams in unseen language models based on whether these n-grams are predicted to decrease or increase tracking error. We show that this algorithm can reduce the tracking error rate by more than 24% on unseen test data.