

# 2007 Student Research Symposium

*Language Technologies Institute*



## Fast learning of Document Ranking Functions with the Committee Perceptron

Jonathan Elsas

### **Abstract:**

This paper presents a new variant of the perceptron algorithm using selective committee averaging (or voting). We apply this algorithm to the problem of learning ranking functions for document retrieval, known as the "Learning to Rank" problem. Most previous algorithms proposed to address this problem focus on minimizing the number of mis-ranked document pairs in the training set. The committee perceptron algorithm improves upon existing solutions by biasing the final solution towards maximizing an arbitrary rank-based performance metrics. This method performs comparably or better than two state-of-the-art rank learning algorithms, and also provides significant training time improvements over those methods, showing over a 45-fold reduction in training time compared to ranking SVM.