



Dissertation Defense


Laura Burdick

Understanding Word Embedding Stability Across Languages and Applications

Wednesday, July 15, 2020

2:00 – 4:00 pm

umich.zoom.us/j/97668320384



ABSTRACT: Despite the recent popularity of word embedding methods, there is only a small body of work exploring the limitations of these representations. In this thesis, we consider several aspects of embedding spaces, including their stability. First, we propose a definition of stability, and show that common English word embeddings are surprisingly unstable. We explore how properties of data, words, and algorithms relate to instability. We extend this work to approximately 100 world languages, considering how linguistic typology relates to stability. Additionally, we consider contextualized output embedding spaces. Using paraphrases, we explore properties and assumptions of BERT, a popular embedding algorithm.

Second, we consider how stability and other word embedding properties affect tasks where embeddings are commonly used. We consider both word embeddings used as features in downstream applications, and corpus-centered applications, where embeddings are used to study characteristics of language and individual writers. In addition to stability, we also consider other word embedding properties, specifically batching and curriculum learning, and how methodological choices made for these properties affect downstream tasks.

Finally, we consider how knowledge of stability affects how we use word embeddings.

Throughout this thesis, we discuss strategies to mitigate instability and provide analyses highlighting the strengths and weaknesses of word embeddings in different scenarios and languages. We show areas where more work is needed to improve embeddings, and we show where embeddings are already a strong tool.

Chair: Prof. Rada Mihalcea