

Exploring Lexical Relations in BERT using Semantic Priming

Kanishka Misra

Purdue University, West Lafayette, Indiana, United States

Allyson Ettinger

University of Chicago, Chicago, Illinois, United States

Julia Rayz

Purdue University, West Lafayette, Indiana, United States

Abstract

BERT is a language processing model trained for word prediction in context, which has shown impressive performance in natural language processing tasks. However, the principles underlying BERT's use of linguistic cues present in context are yet to be fully understood. In this work, we develop tests informed by the semantic priming paradigm to investigate BERT's handling of lexical relations to complete a cloze task (Taylor, 1953). We define priming to be an increase in BERT's expectation for a target word (pilot), in a context (e.g., I want to be a ___), when the context is prepended with a related word (airplane) as opposed to an unrelated one (table). We explore BERT's priming behavior under various predictive constraints placed on the blank, and find that BERT is sensitive to lexical priming effects only under minimal constraint from the input context. This pattern was found to be consistent across diverse lexical relations.