

Compositional Neural Machine Translation by Removing the Lexicon from Syntax

Tristan Thrush

Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Abstract

The meaning of a natural language utterance is largely determined from its syntax and words. Additionally, there is evidence from theories in semantics and neuroscience that humans process an utterance by separating some amount of knowledge about the lexicon from the knowledge of word order. In this paper, we propose neural units that can enforce this constraint over an LSTM encoder and decoder. We demonstrate that our model achieves competitive performance across a variety of domains including semantic parsing, syntactic parsing, and English to Mandarin Chinese translation. In these cases, our model outperforms the standard LSTM encoder and decoder architecture on many or all of our metrics. To demonstrate that our model achieves a desired partial separation between the lexicon and syntax, we analyze its weights and explore its behavior when different neural modules are damaged. When damaged, we find that the model displays the knowledge distortions that aphasics are evidenced to have.