




Zero-shot Neural Transfer for Cross-lingual Entity Linking



Shruti Rijhwani
Carnegie Mellon
University


Cross-lingual entity linking maps a named entity in a source language to its corresponding entry in a structured knowledge base that is in a different (target) language. While previous work relies heavily on bilingual lexical resources to bridge the gap between the source and the target languages, these resources are scarce or unavailable for many low-resource languages. To address this problem, we investigate zero-shot cross-lingual entity linking, in which we assume no bilingual lexical resources are available in the source low-resource language. Specifically, we propose pivot-based entity linking, which leverages information from a high-resource "pivot" language to train character-level neural entity linking models that are transferred to the source low-resource language in a zero-shot manner. With experiments on nine low-resource languages and transfer through a total of 54 languages, we show that our proposed pivot-based framework improves entity linking accuracy 17% (absolute) on average over the baseline systems for the zero-shot scenario. Further, we also investigate the use of language-universal phonological representations which improves average accuracy (absolute) by 36% when transferring between languages that use different scripts.

Cross-lingual Alignment vs Joint Training:

A Comparative Study and A Simple Unified Framework

Learning multilingual representations of text has proven a successful method for many cross-lingual transfer learning tasks. There are two main paradigms for learning such representations: (1) alignment, which maps different independently trained monolingual representations into a shared space, and (2) joint training, which directly learns unified multilingual representations using monolingual and cross-lingual objectives jointly. In this work, we first conduct direct comparisons of representations learned using both of these methods across diverse cross-lingual tasks. Our empirical results reveal a set of pros and cons for both methods, and show that the relative performance of alignment versus joint training is task-dependent. Stemming from this analysis, we propose a simple and novel framework that combines these two previously mutually-exclusive approaches. We show that our proposed framework alleviates limitations of both approaches and can generalize to contextualized representations such as Multilingual BERT.

Zirui Wang
Carnegie Mellon
University



Friday, September 25, 2020

1:30 - 2:50 PM EST

Join the meeting on Zoom

Meeting ID 968 6753 2227 Passcode 883155