

AN EMPIRICAL STUDY ON FINDING SPANS

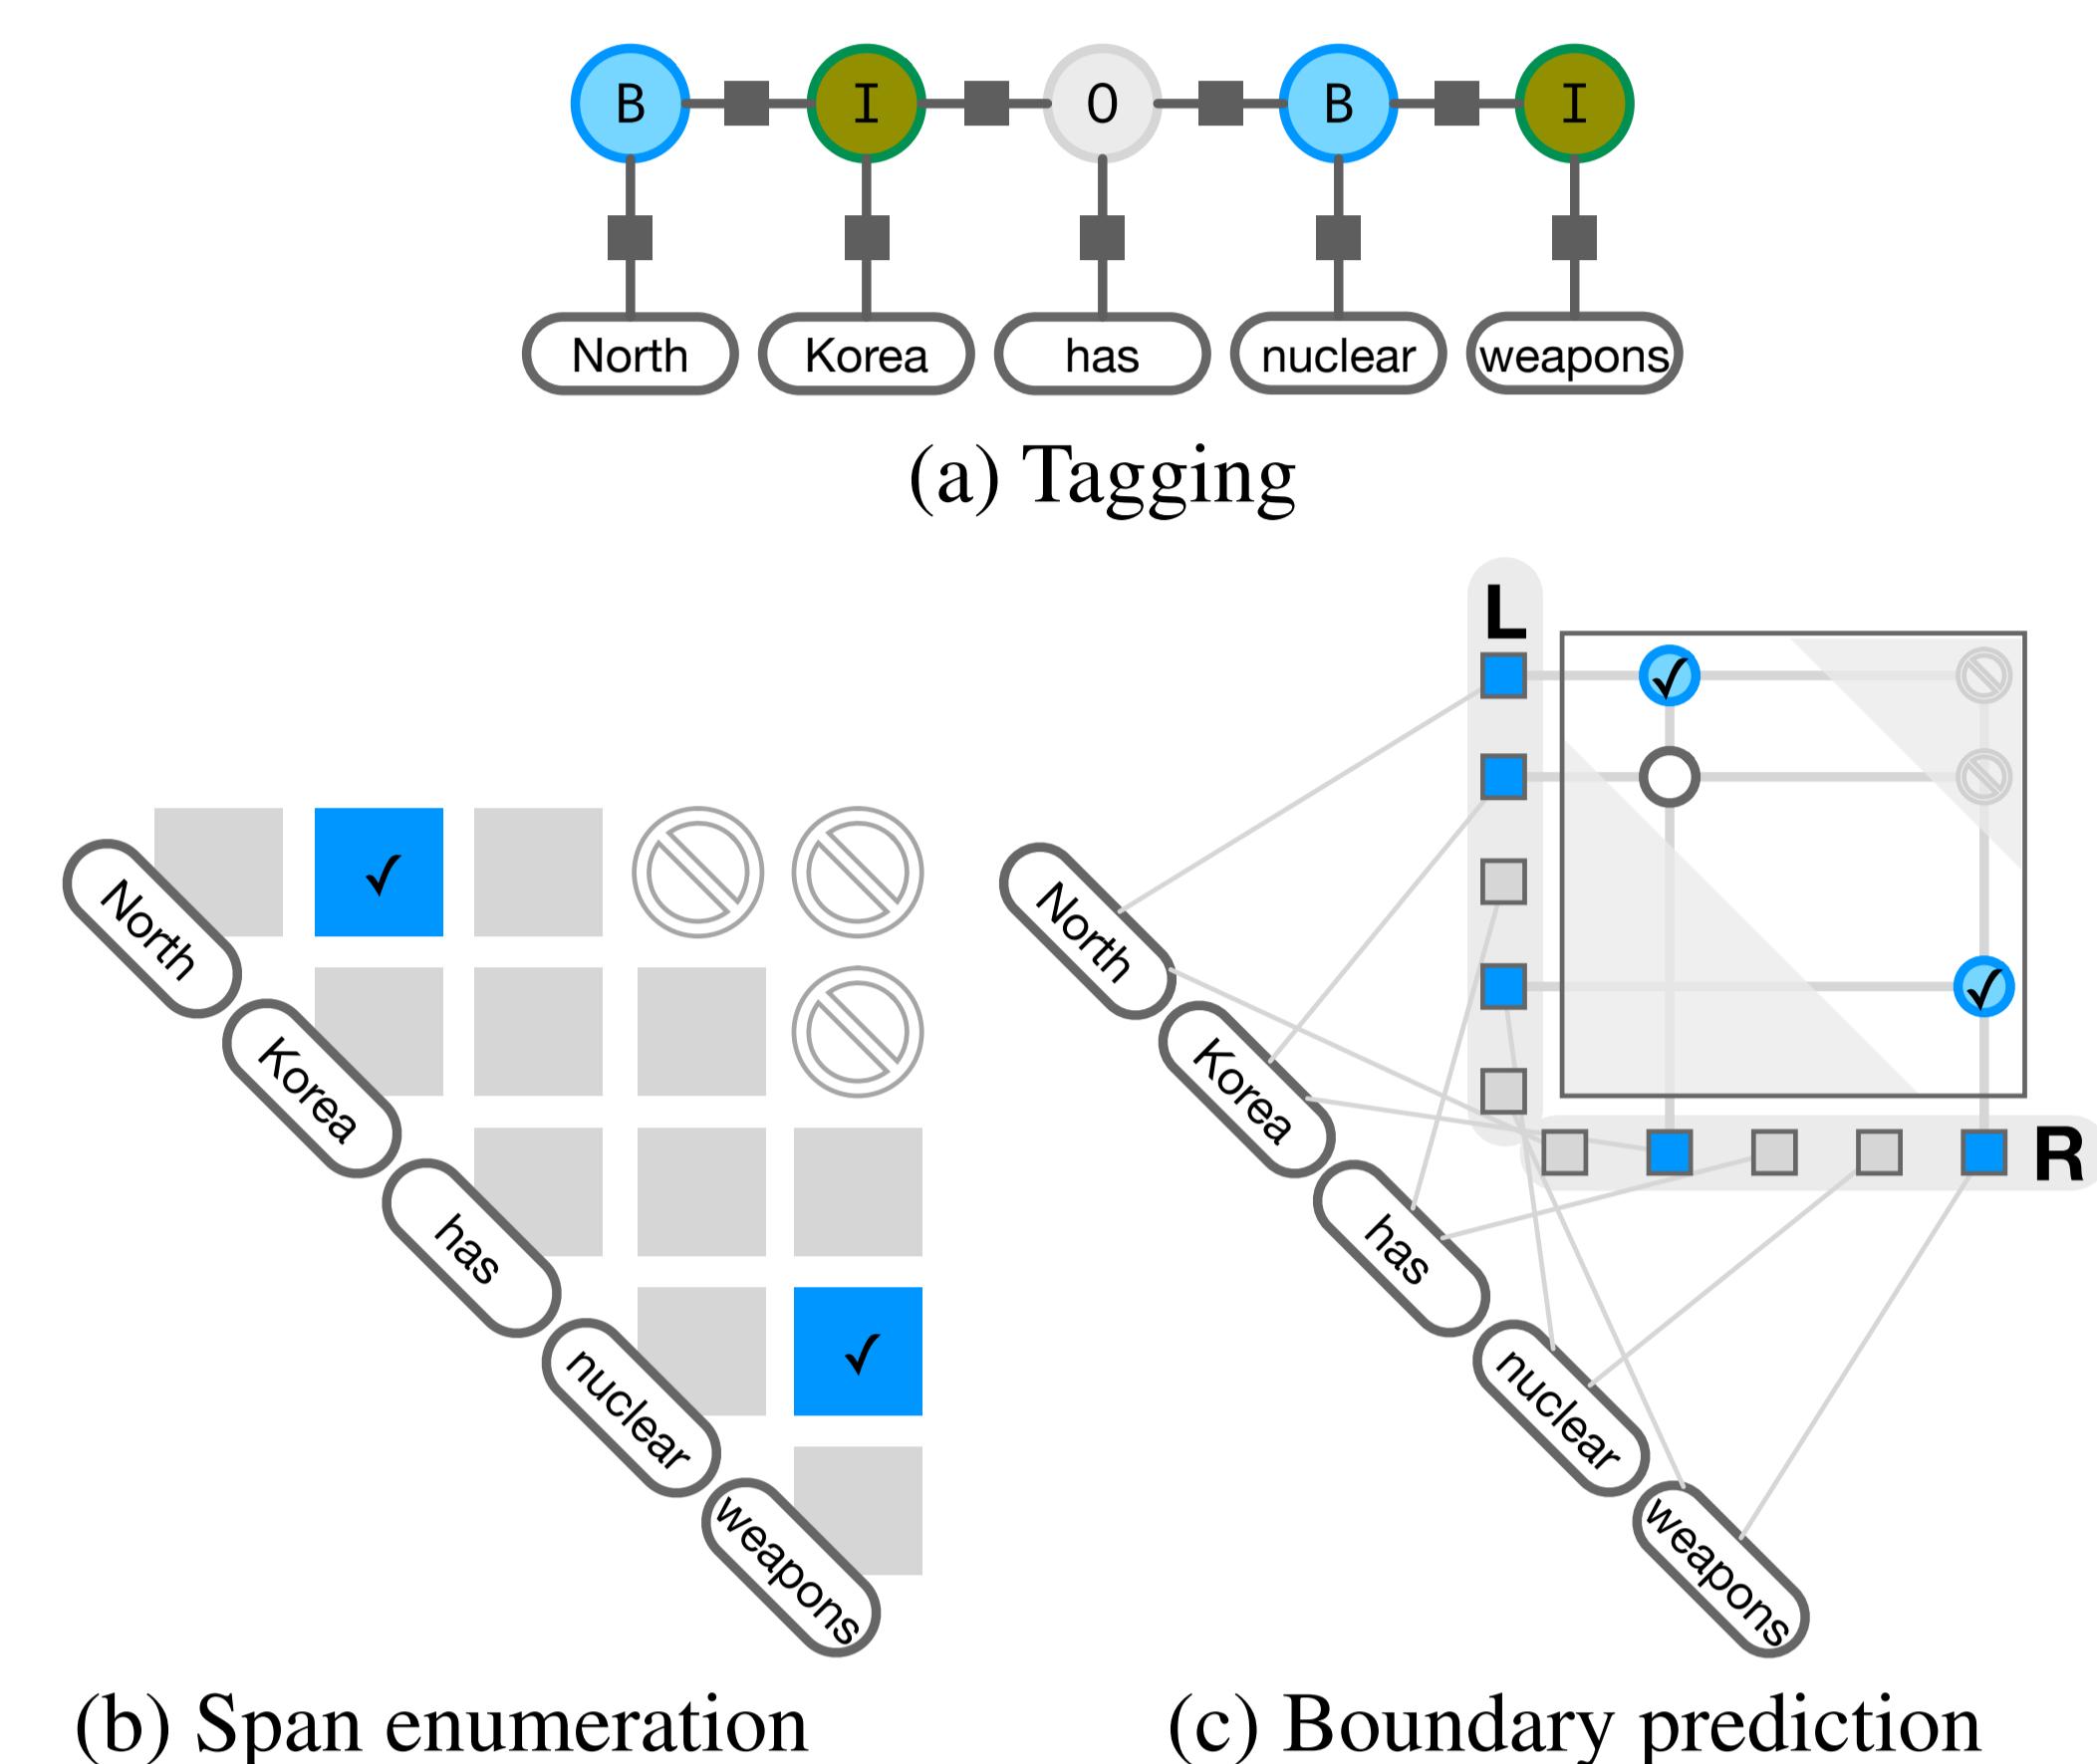
Weiwei Gu^{*1}, Boyuan Zheng^{*2}, Yunmo Chen², Tongfei Chen³, Benjamin Van Durme²

¹University of Rochester, ²Johns Hopkins University, ³Microsoft Semantic Machines

Summary

- The best span finding method depends on downstream tasks:
 - Tagging has higher precision
 - Span enumeration and boundary prediction have higher recall
 - Boundary prediction is preferable to span enumeration for most cases
- The inclusion of mention type helps on boundary prediction, but not on tagging or span enumeration
- An additional contextualized LSTM-RNN layer might be helpful for T5 encoder (an encoder-decoder model), but not on RoBERTa (encoder only model).

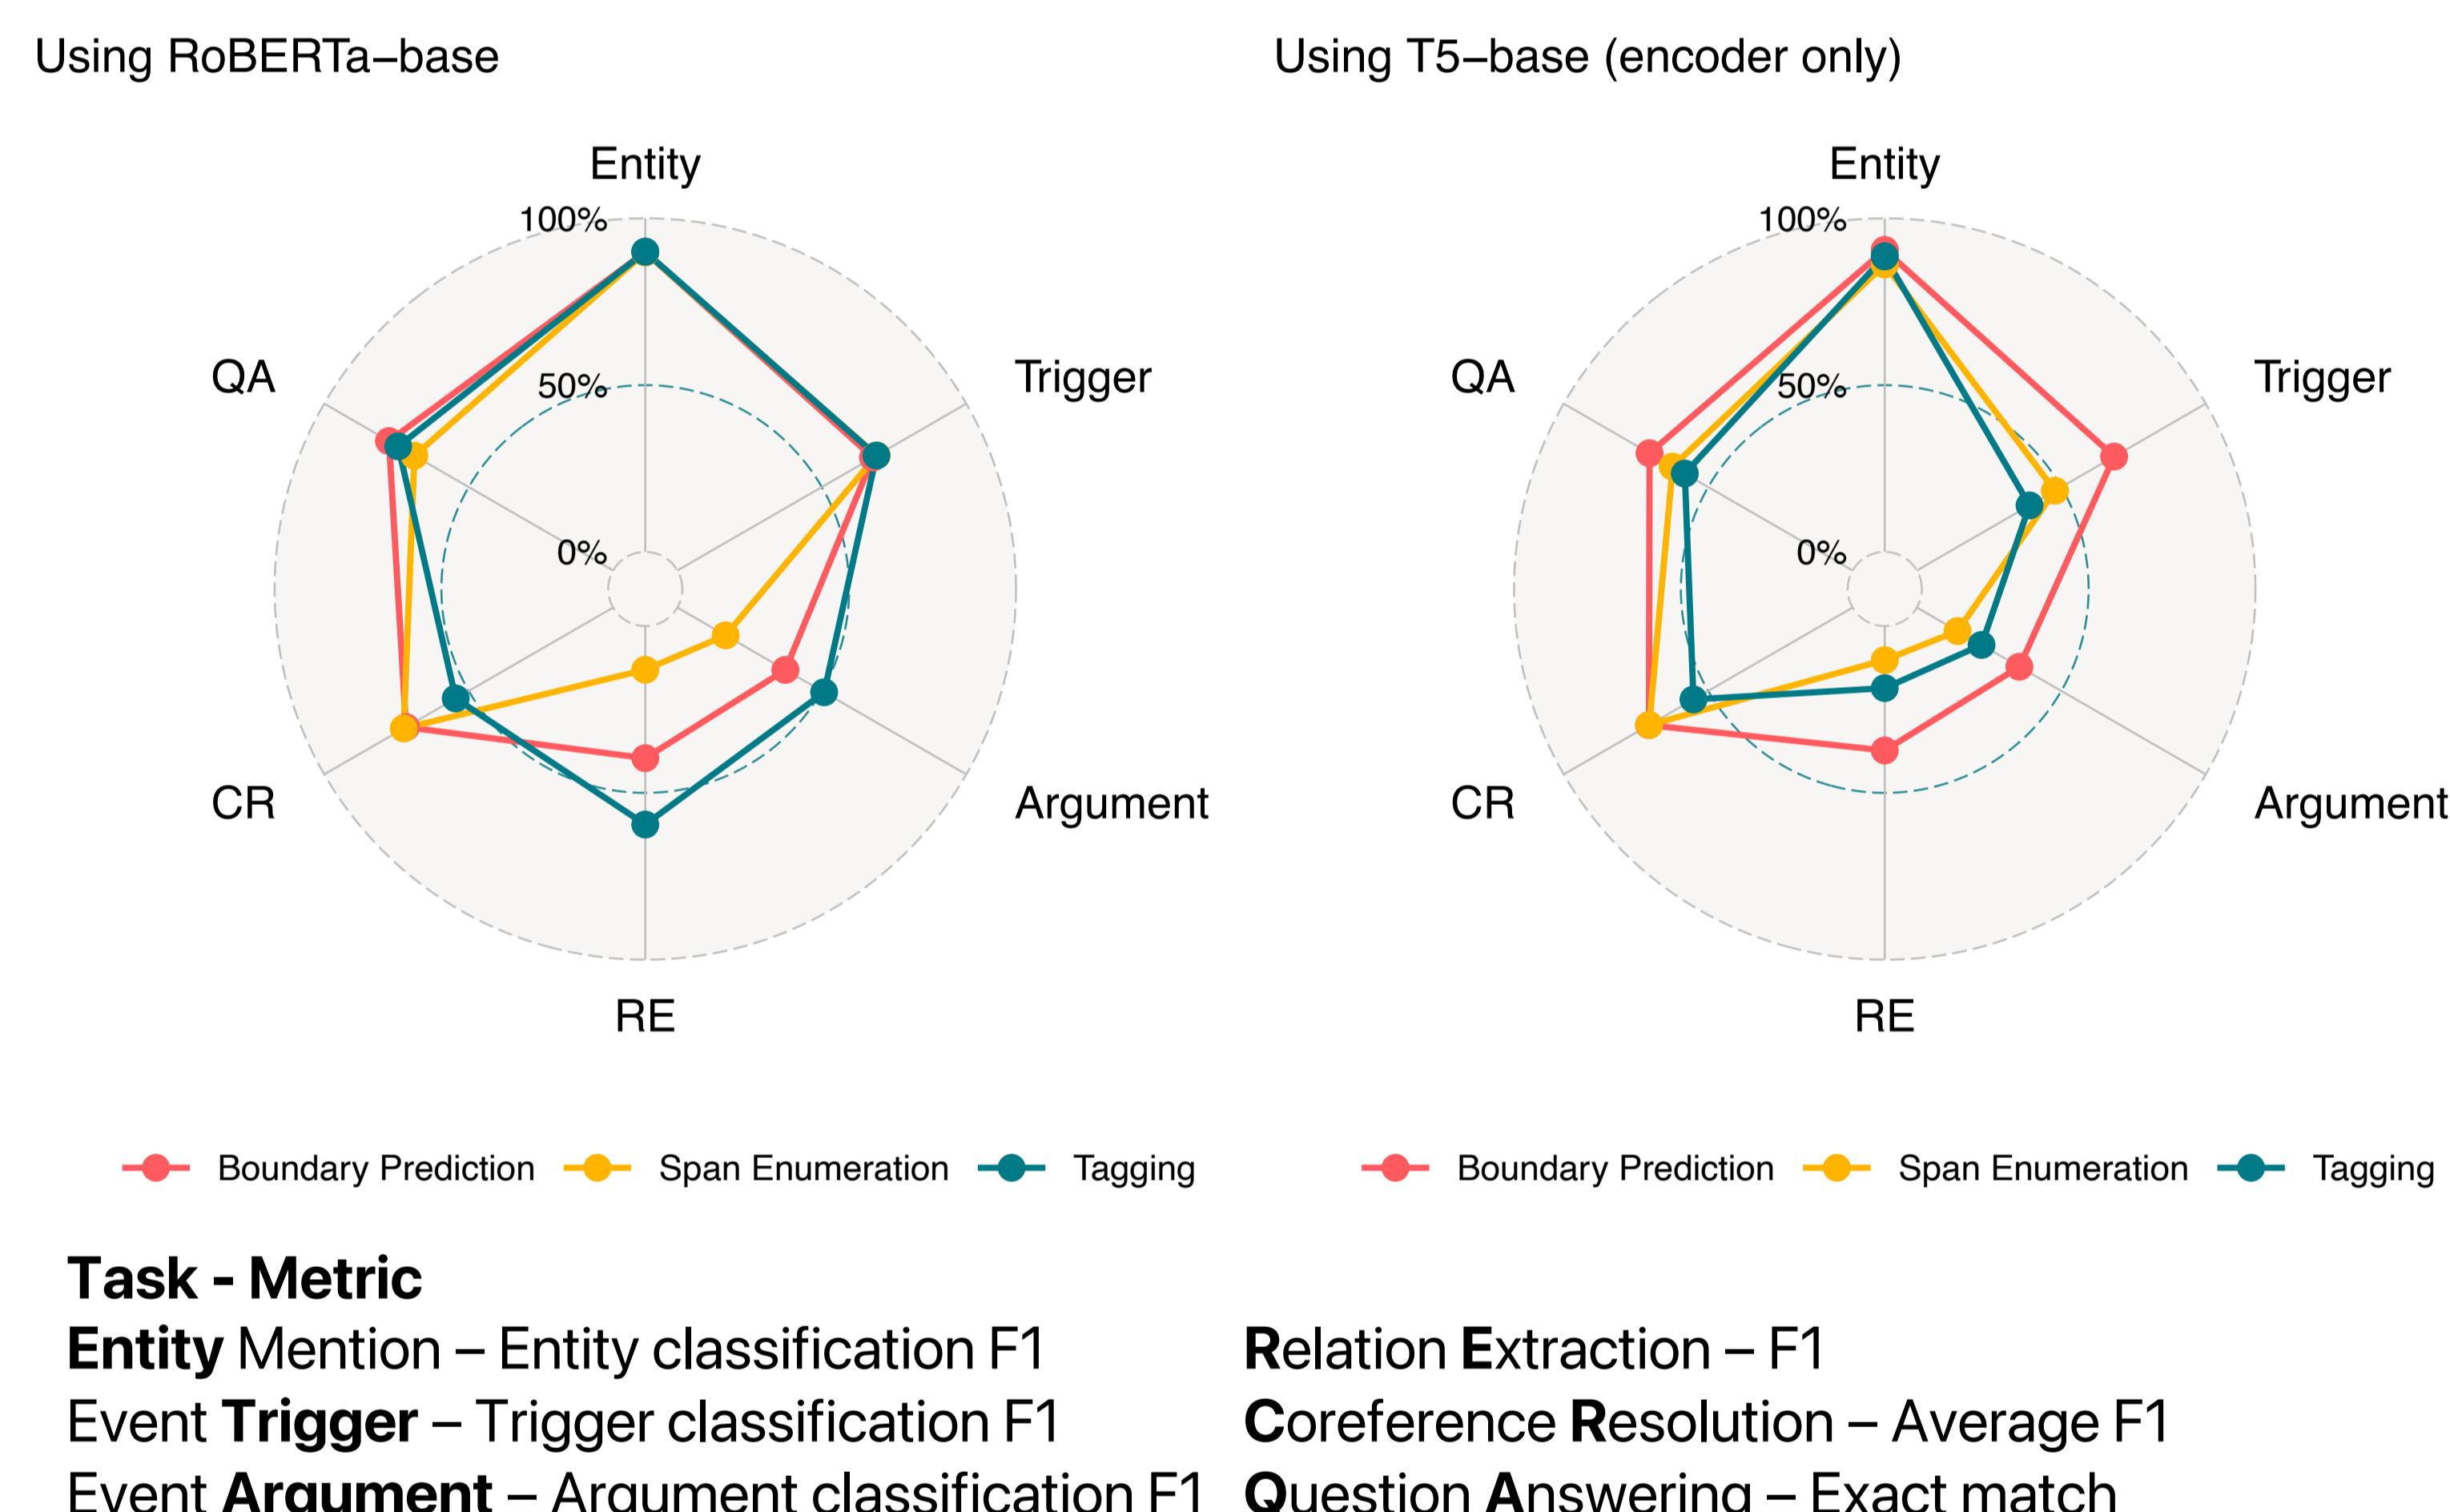
Different Span Finding Methods



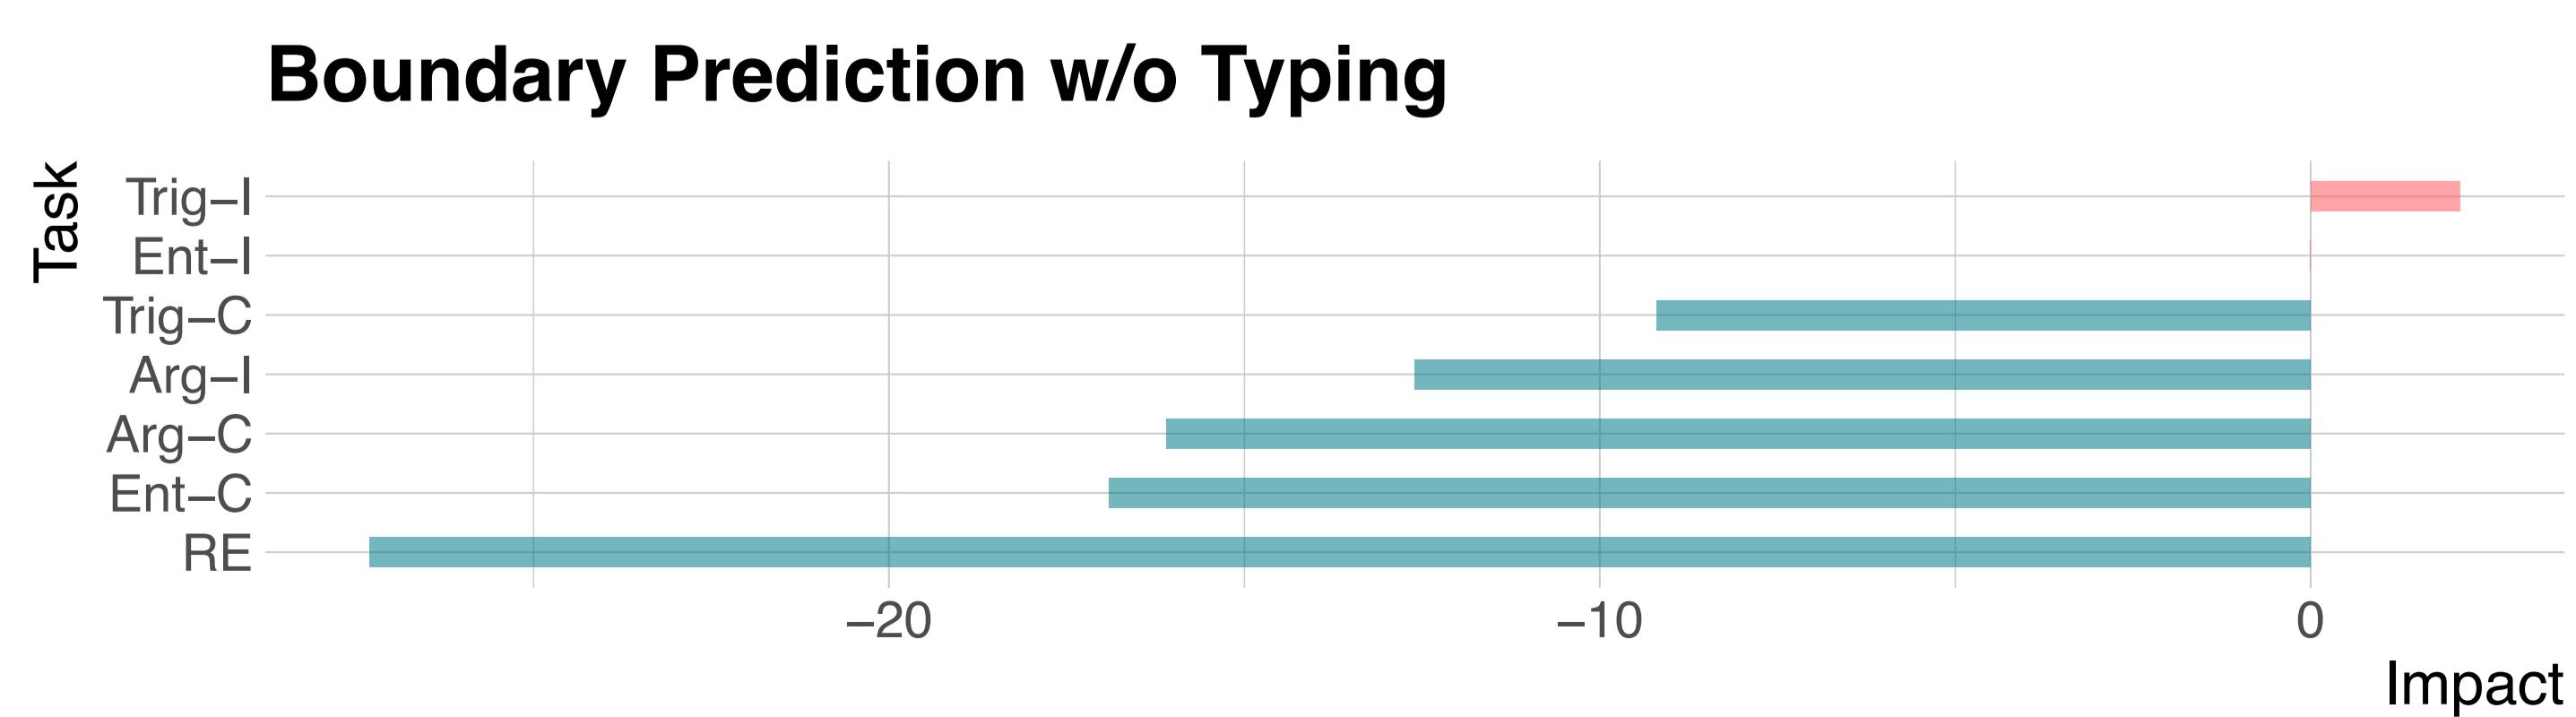
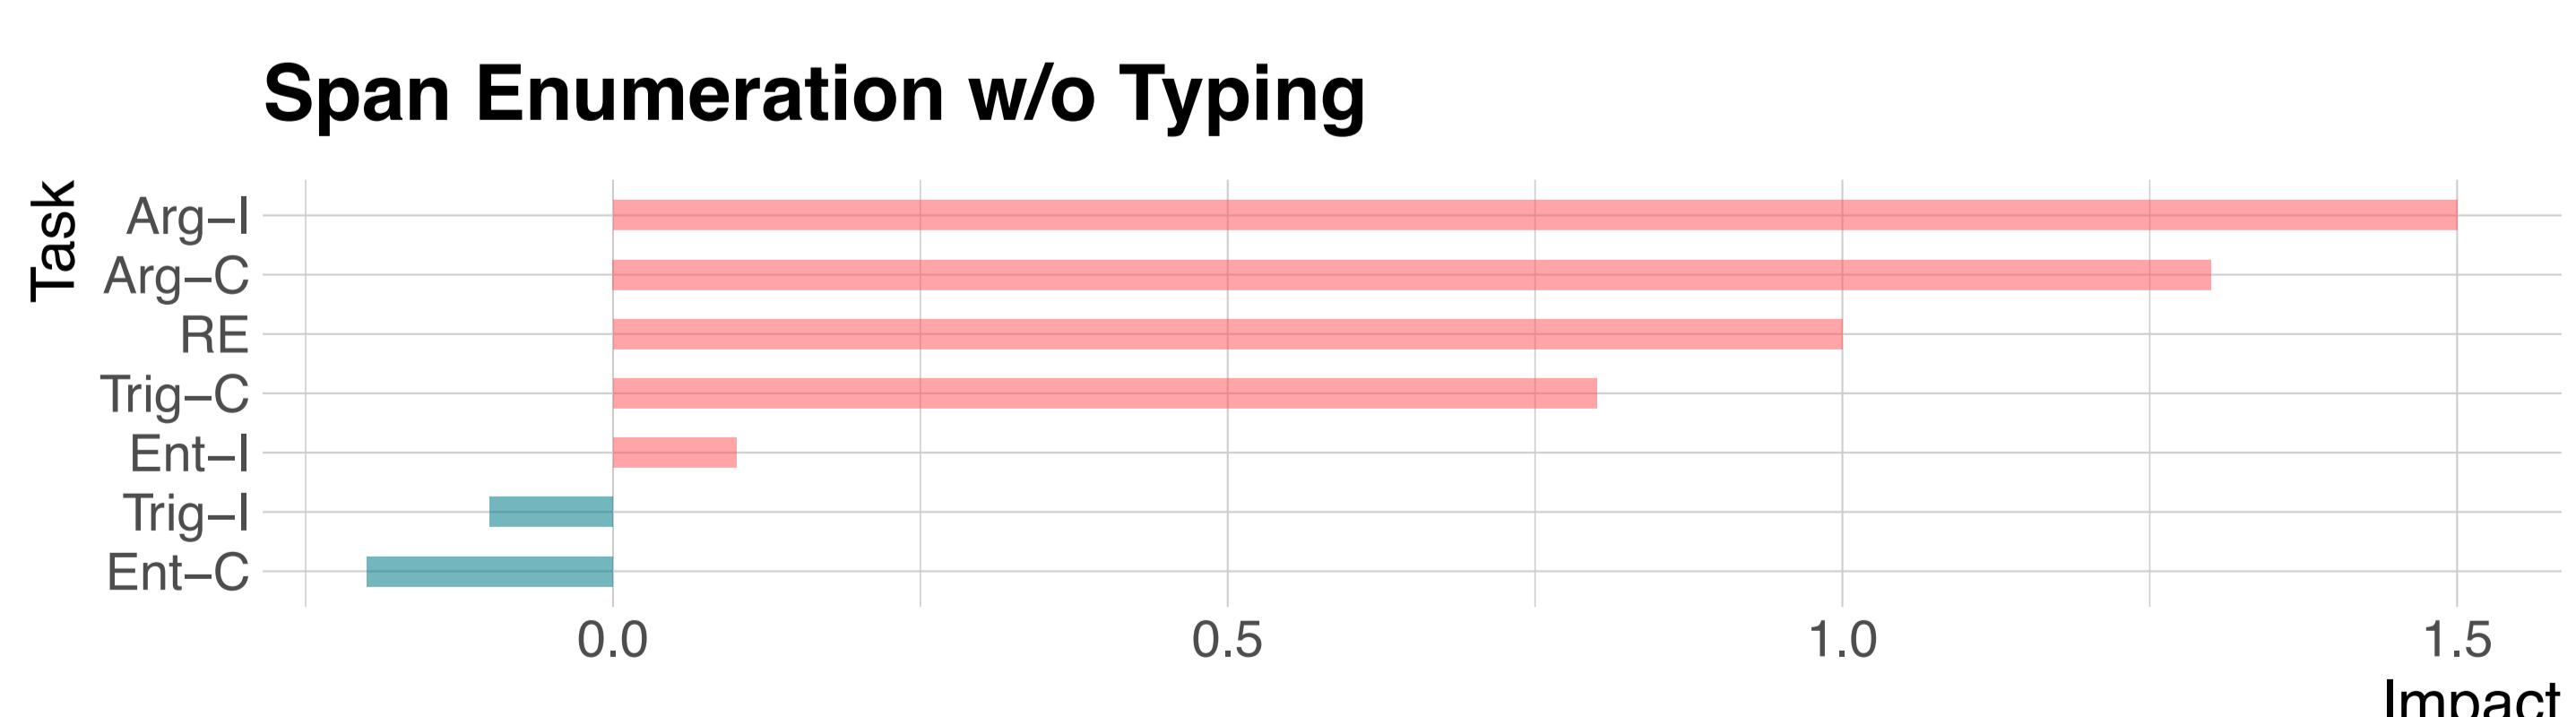
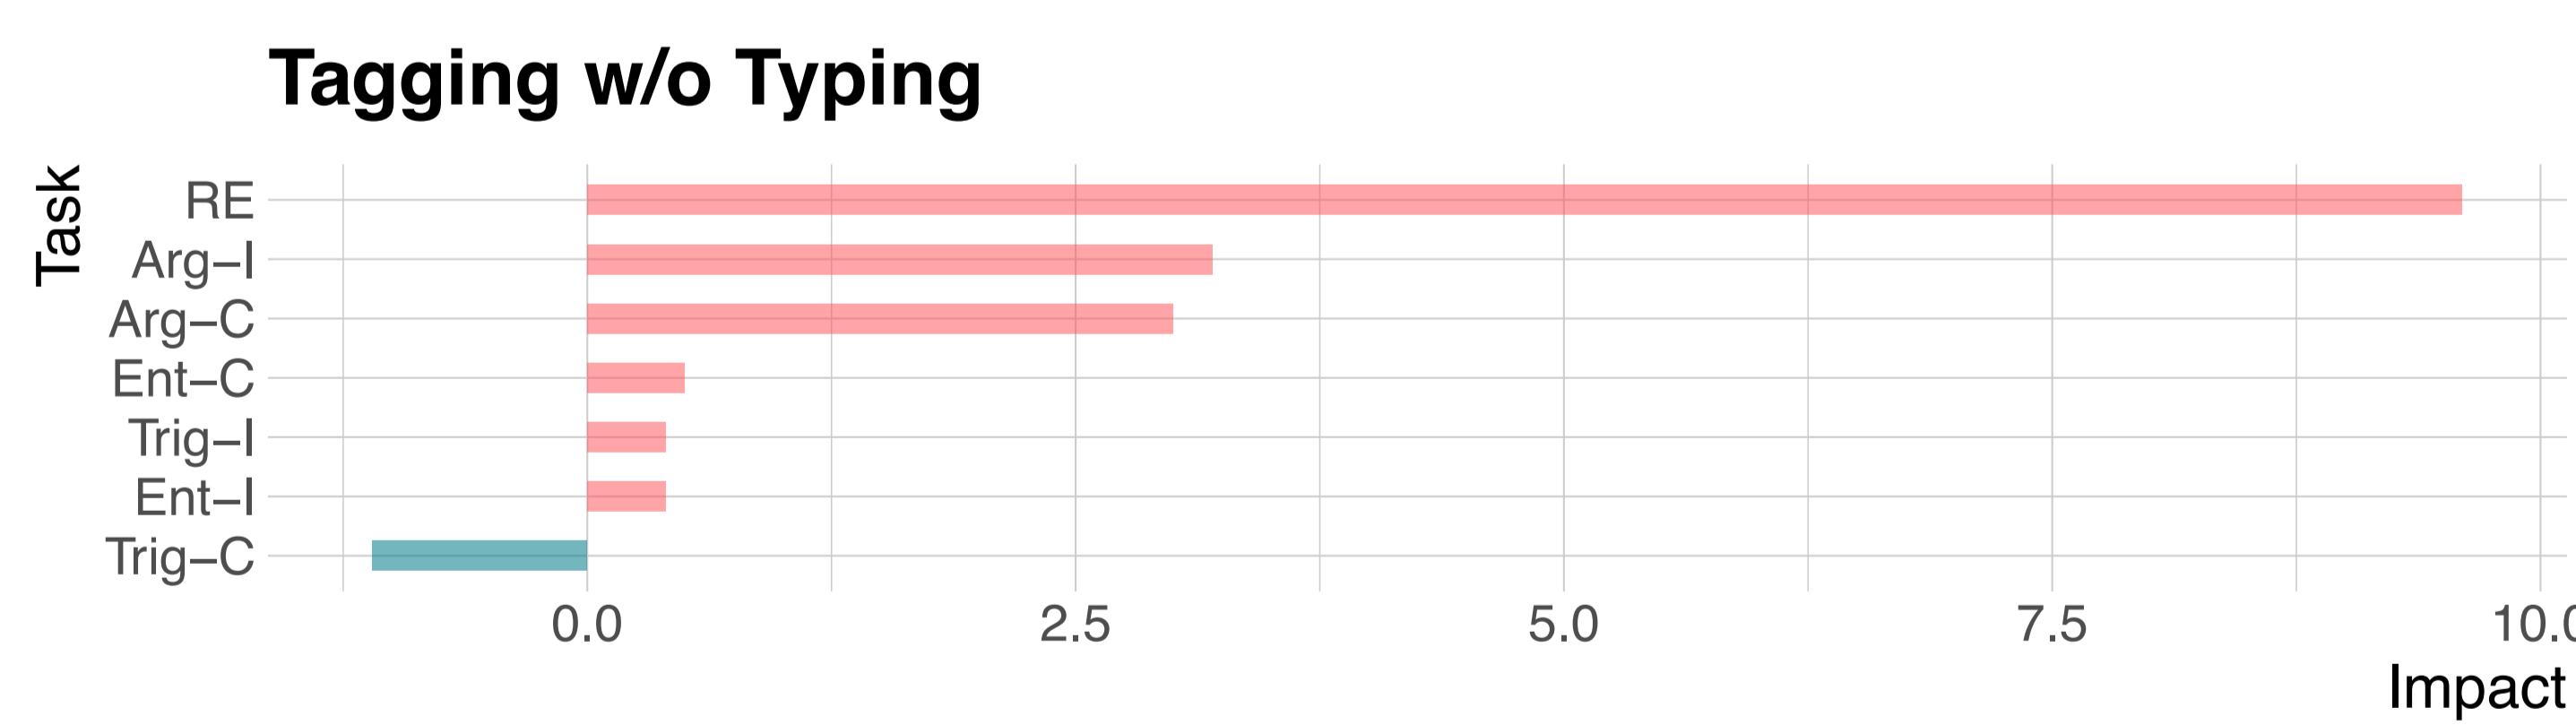
Experiment Setup

- NER, EE, and RE:
 - Dataset: ACE05
 - Model: OneIE (Lin et al., 2020) with global features disabled
- Coreference Resolution(CR):
 - Dataset: OntoNotes
 - Model: Higher-order Coreference Resolution (Lee et al., 2018)
- Extractive QA:
 - Dataset: SQuAD2.0 Validation set
 - Model: Transformer QA (Devlin et al., 2019)
- Encoders: RoBERTa-base and T5-base (encoder-only)

Which Finder For Your Task?



Typing Information?



Additional Contextualization?

