

Relationship extraction is the task of extracting semantic relationships between entities from a text. We create a Czech Relationship Extraction Dataset (CERED) using distant supervision on Wikidata and Czech Wikipedia. We detail the methodology we used and the pitfalls we encountered. Then we use CERED to fine-tune a neural network model for relationship extraction. We base our model on BERT – a linguistic model pre-trained on extensive unlabeled data. We demonstrate that our model performs well on existing English relationship datasets (Semeval 2010 Task 8, TACRED) and report the results we achieved on CERED.