Anaphora Resolution via Latent Antecedents Representation: A Brief Idea

Matteo Grella
matteogrella@gmail.com

May 13, 2018

Abstract

An idea is proposed to perform the task of anaphora resolution as a pointing problem (Lee et al., 2017) using the bidirectional recurrent autoencoders schema introduced by Grella and Cangialosi (2018).

Anaphora resolution is a core task in natural language processing. It consists to resolve what a pronoun or a noun refers to by finding its antecedent i.e. another linguistic expression that shares the same meaning in a discourse.

I propose a neural-based approach to perform anaphora resolution via semi-supervised learning that relies on the ability of the autoencoders (Rumelhart et al., 1985) to learn a representation with the purpose of reconstructing its own input, and in particular on the **bidirectional recurrent autoencoders** (Grella and Cangialosi, 2018) capable to reconstruct for each *i-th* input another *j-th* input of the same sequence. The resulting approximate representation (or "latent representation") can be then converted to a position in the input sequence using a simple vector similarity.

The model is composed of two BiLSTM. The first BiLSTM receives in input the tokens of a sentence already encoded in a dense "meaningful" representation (e.g. pre-trained word and pos embeddings). It encodes the tokens into *context vectors* that represent them with their surrounding context.

The *context vectors* are in turn given to the second BiLSTM that acts as an autoencoder, transforming them into *latent antecedents*. It is trained to reconstruct each *context vector* with the one that represents its antecedent, if it exists, or with itself if not, by minimizing the mean absolute error.

The errors are propagated all the way back, through the two BiLSTM until the initial tokens embeddings. An optimizer is used to update the parameters according to the gradients.

Considering that the network learns its own suitable representation for both the input and the output by itself and that we give a "teaching signal" only (that is the index of the *context vector* of the antecedent to reconstruct without imposing any particular representation), I think that it is possible to consider the approach a semi-supervised one.

During the training, the antecedents are taken from a dataset. At test time, a vector similarity is used to point to the antecedents, comparing the *latent antecedents* with the *context vectors*.

The proposed model should automatically capture the syntactic and semantic features relevant to the task at hand. I plan to evaluate the contribution resulting from the integration of this information into a neural dependency parsing model, and vice versa to test the impact of external syntactic knowledge for the anaphora resolution.