Adverse Reaction Extraction from Drug Labels Using LSTM Networks

Xuan Zhu, Mengxin Qian, David Guy Brizan
Master of Science in Analytics, University of San Francisco
San Francisco, CA

System Overview

Our approach is to apply Deep Learning models (specifically, Long Short-Term Memory (LSTM) networks) on trained word embedding to extract AdverseReactions and related mentions (Severity, Factor, DrugClass, Negation, Animal).

Data Description

Training data of drug labels was provided by FDA in the XML format via the TCA 2017 website. Additional unannotated XML data was used for submission and evaluation.

Model Description

LSTM is a type of Recurrent Neural Network (RNN) that learn how to "forget" past observations. This makes them more robust to noise, and better able to capture “long-term dependencies” in a sequence. Sentences are made of the sequence of words and order of words encode a lot of information that is useful to predict sentiment. Thus, the essential idea is to adopt the LSTM architecture to capture something like “longer-than-bigram” dependencies in sentences. This approach can be described in two major steps: 1) map words to word embeddings (e.g. train an embedding with word2vec or Keras); 2) train the RNN with a sequence of vectors as input and considers the order of the vectors to generate a prediction.

In the second step, the embedding layer (the new representations of sentences) will be passed to LSTM cells. These will add recurrent connections to the network so we can include information about the sequence of words in the data. Finally, the LSTM cells will go to a softmax output layer to give a single unit output, because we are trying to predict if this word is an entity of AdverseReactions, Severity, Factor, DrugClass, Negation, Animal, or none of the above.

Results

Such LSTM-based entity classification system demonstrated high validation accuracy if provided with proper labelled training dataset. However, our results also indicated that such models usually require significant efforts on fine-tuning and is challenge to assess and prevent overfitting. Next steps may include adding more language (NLP) features from natural language processing approaches to blend NLP and advanced learning for domain-specific name entity recognition problems.