

Deep Learning in NLP: Summary / Takeaway

Christian Wurm & Tatiana Bladier

Heinrich-Heine-Universität Düsseldorf
Wintersemester 2018/2019

Overview

- 1 Course contents
- 2 Further exploration of DL in NLP
- 3 AP Projects

What was it all about?

- ★ Introductory course to Deep Learning in NLP and applications with Keras.
- ★ Now you are able to understand and build neural network models using Keras and Python.
- ★ You can solve some NLP Problems using deep learning architectures: POS-tagging, named entity recognition, chunking, sentiment analysis, text classification, semantic role labeling, text generation.
- ★ You can better understand and adapt code you find on GitHub according to your deep learning task (or write your own code based on some scientific article).

Deep Learning Projects: 6 Basic Steps

- 1 Problem or task: machine translation, sentiment analysis, text classification, part-of-speech tagging, text generation etc.
- 2 Input data points (features): files of people speaking, text files, images etc.
- 3 Examples of the expected outputs (tags or labels): transcripts of sound files, positive/negative, cat/dog/fish etc.
- 4 Suitable algorithm to train on those data.
- 5 Way to evaluate how good the algorithm and the model is (distance between the algorithm's output and expected output).
- 6 Present your results.

Neural Network Models we used in class

- ★ Feed-Forward Network (FFN) (Boston housing)
- ★ Simple Recurrent Neural Network (RNN) (Name recognition)
- ★ Long Short-Term Memory (LSTM) (POS-tagging)
- ★ Gated Recurrent Unit (GRU) + Bidirectional GRU (used for sequence labeling)
- ★ Bidirectional Recurrent Neural Network (BiLSTM) (POS-tagging, chunking, named entity recognition)
- ★ Convolutional Neural Network (CNN) for character models
→ capture morphological information
- ★ Multichannel CNN (sentiment analysis)
- ★ Conditional Random Field (CRF) + BiLSTM (semantic role labeling, sequence labeling in general)

Use embeddings

- ★ Character-level embeddings (to encode characters in words).
- ★ Word-level embeddings (one-hot, pretrained, own embeddings: Word2Vec).
- ★ Sentence-level embeddings + text-level embeddings.

What else?

- ★ Use Keras and Keras functional API.
- ★ Prepare (vectorize) your data for training.
- ★ Visualize your neural network model (using plt model).
- ★ Save and re-use pre-trained models.
- ★ Make predictions with the model.
- ★ Evaluate the model performance with the following metrics: accuracy, recall + precision + F1 score, confusion matrix.
- ★ Use early stopping to save the best model and stop training.
- ★ Use fit generator to avoid out-of-memory (OOM) issues.
- ★ Use LaTeX to describe your experiments and results.

Books on Deep Learning in NLP

- ★ Chollet (2017) *Deep Learning with Python*.
- ★ Goldberg and Hirst (2017) *Neural Network Methods for Natural Language Processing*.
- ★ Brownlee (2016). *Deep Learning for Natural Language Processing. Develop Deep Learning Models for Natural Language in Python*.
- ★ Geron, Aurelien (2017) *Hands-on machine learning with Scikit-Learn and TensorFlow : concepts, tools, and techniques to build intelligent systems*.
- ★ Harrison and Honnibal (2020) *Deep Learning with Text: A Modern Approach to Natural Language Processing with Python and Keras*.

Some good tutorials

- ★ <https://keras.io/>
- ★ <https://machinelearningmastery.com/>

AP Projects: Evaluation Criteria

- ★ Clarity of writing / presentation: 1-5
- ★ Research question / hypothesis: 1-5
- ★ Soundness / correctness: 1-5
- ★ Empirical evaluation: 1-5
- ★ Dataset / effort: 1-5

References I

- Brownlee, J. (2016). *Deep Learning for Natural Language Processing. Develop Deep Learning Models for Natural Language in Python.*
- Chollet, F. (2017). *Deep Learning with Python.* Manning Publications.
- Geron, Aurelien (2017). *Hands-on machine learning with Scikit-Learn and TensorFlow : concepts, tools, and techniques to build intelligent systems.* O'Reilly Media, Sebastopol, CA.
- Goldberg, Y. and Hirst, G. (2017). *Neural Network Methods for Natural Language Processing.* Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- Harrison, P. and Honnibal, M. (2020). *Deep Learning with Text: A Modern Approach to Natural Language Processing with Python and Keras.* O'Reilly Media.