

# Notes on NLP

## Papers / Websites

- The Annotated Transformer (Harvard): <http://nlp.seas.harvard.edu/2018/04/03/attention.html>
- Attention Is All You Need (original paper): <https://arxiv.org/abs/1706.03762>
- Distilling the Knowledge in a Neural Network: <https://arxiv.org/abs/1503.02531>
- Well-Read Students Learn Better: On the Importance of Pre-training Compact Models: <https://arxiv.org/abs/1908.08962>
- TensorFlow Hub: <https://tfhub.dev/>
- Google Research, BERT - Smaller Models on Git: <https://github.com/google-research/bert>
- BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding: <https://arxiv.org/pdf/1810.04805.pdf>

## Websites

-  Getting meaning from text: self-attention step-by-step video (Romain Futrzynski): <https://peltarion.com/blog/data-science/self-attention-video>
-  The Illustrated Transformer (Jay Alammar): <http://jalammar.github.io/illustrated-transformer/>
-  Paper Dissected: "Attention is All You Need" Explained: <https://mlexplained.com/2017/12/29/attention-is-all-you-need-explained/>
- Speech and Language Processing, Dan Jurafsky and James H. Martin (3rd ed. draft): <https://web.stanford.edu/~jurafsky/slp3/>

## Videos

- RNN W3L08 - Attention Model (Andrew Ng): <https://www.youtube.com/watch?v=FMXUkEbjf9k&feature=youtu.be>
- Attention is all you need; Attentional Neural Network Models | Łukasz Kaiser | Masterclass: <https://www.youtube.com/watch?v=rBCqOTefxvg&feature=youtu.be>
- Transformer Neural Networks - EXPLAINED! (Attention is all you need) (CodeEmporium): <https://www.youtube.com/watch?v=TQQIZhbC5ps&feature=youtu.be>
- Attention in Neural Networks (CodeEmporium): <https://www.youtube.com/watch?v=W2rWgXJBZhU&t>
-  [Transformer] Attention Is All You Need | AISC Foundational (Joseph Palermo (Dessa)): [https://www.youtube.com/watch?v=S0KakHcj\\_rs&feature=youtu.be](https://www.youtube.com/watch?v=S0KakHcj_rs&feature=youtu.be)
- Ivan Bilan: Understanding and Applying Self-Attention for NLP | PyData Berlin 2018: <https://www.youtube.com/watch?v=OYyqPG4d9H0&feature=youtu.be>

-  Transformer (Attention is all you need)(Minsuk Heo): <https://www.youtube.com/watch?v=z1xs9jdZnuY&feature=youtu.be>
-  Self-attention step-by-step | How to get meaning from text | Peltarion Platform (Romain Futrzynski): [https://www.youtube.com/watch?v=-9vVhYEXeyQ&feature=emb\\_logo](https://www.youtube.com/watch?v=-9vVhYEXeyQ&feature=emb_logo)

# Literature overview on NLP

These tables should give an overview over recent and influential literature in the field of Natural Language Processing from the past few years.

## General overview

NLP, transfer learning, language models.

Author	Title	Link to code	Abstract (short)
Vaswani et al. (2017)	<a href="#">Attention Is All You Need</a>	Code used for training and evaluation: <a href="https://github.com/tensorflow/tensor2tensor">https://github.com/tensorflow/tensor2tensor</a>	Introduction of a new simple network architecture, the <b>Transformer</b> , based solely on attention mechanisms, dispensing with recurrence and convolutions entirely.
Kim et al. (2017)	<a href="#">Structured Attention Networks</a>	<a href="https://github.com/harvardnlp/struct-attn">https://github.com/harvardnlp/struct-attn</a>	In this work, we experiment with incorporating richer structural distributions, encoded using graphical models, within deep networks. We show that these structured <b>attention</b> networks are simple extensions of the basic attention procedure, and that they allow for extending attention beyond the standard soft-selection approach, such as attending to partial segmentations or to subtrees.

Author	Title	Link to code	Abstract (short)
Radford et al. (2018)	<a href="#">Improving Language Understanding by Generative Pre-Training</a>	<a href="https://github.com/openai/finetune-transformer-lm">https://github.com/openai/finetune-transformer-lm</a>	<p>Natural language understanding comprises a wide range of diverse tasks such as textual entailment, question answering, semantic similarity assessment, and document classification. Although large unlabeled text corpora are abundant, labeled data for learning these specific tasks is scarce, making it challenging for discriminatively trained models to perform adequately. We demonstrate that large gains on these tasks can be realized by generative pre-training of a language model on a diverse corpus of unlabeled text, followed by discriminative fine-tuning on each specific task. <b>GPT-1</b></p>
Devlin et al. (2018)	<a href="#">BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding</a>	<a href="https://github.com/openai/finetune-transformer-lm">https://github.com/openai/finetune-transformer-lm</a>	<p>Introduction of a new language representation model called <b>BERT</b>, which stands for Bidirectional Encoder Representations from Transformers. BERT is designed to pre-train deep bidirectional representations from unlabeled text by jointly conditioning on both left and right context in all layers.</p>

Author	Title	Link to code	Abstract (short)
Radford et al. (2019)	<a href="#">Language Models are Unsupervised Multitask Learners</a>	<a href="https://github.com/openai/gpt-2">https://github.com/openai/gpt-2</a>	Natural language processing tasks, such as question answering, machine translation, reading comprehension, and summarization, are typically approached with supervised learning on task-specific datasets. We demonstrate that language models begin to learn these tasks without any explicit supervision when trained on a new dataset of millions of webpages called WebText.(...) Our largest model, <b>GPT-2</b> , is a 1.5B parameter Transformer that achieves state of the art results on 7 out of 8 tested language modeling datasets in a zero-shot setting but still underfits Web-Text.
Ruder (2019)	<a href="#">Neural Transfer Learning for Natural Language Processing</a>	<a href="https://github.com/sebastianruder">https://github.com/sebastianruder</a>	Multiple novel methods for different <b>transfer learning</b> scenarios were presented and evaluated across a diversity of settings where they outperformed single-task learning as well as competing transfer learning methods.
Kovaleva et al. (2019)	<a href="#">Revealing the Dark Secrets of BERT</a>		BERT-based architectures currently give state-of-the-art performance on many NLP tasks, but little is known about the exact mechanisms that contribute to its success. In the current work, we focus on the <b>interpretation of self-attention</b> , which is one of the fundamental underlying components of BERT.

Author	Title	Link to code	Abstract (short)
Rogers et al. (2020)	<a href="#">A Primer in BERTology: What We Know About How BERT Works</a>		<p>This paper is the first survey of over 150 studies of the popular BERT model. We review the current state of knowledge about <b>how BERT works</b>, what kind of information it learns and how it is represented, common modifications to its training objectives and architecture, the overparameterization issue and approaches to compression.</p>
Brown et al. (2020)	<a href="#">Language Models are Few-Shot Learners</a>	<a href="https://github.com/openai/gpt-3">https://github.com/openai/gpt-3</a>	<p>Demonstration that scaling up language models greatly improves task-agnostic, few-shot performance, sometimes even reaching competitiveness with prior state-of-the-art fine-tuning approaches. Specifically, we train <b>GPT-3</b>, an autoregressive language model with 175 billion parameters, 10x more than any previous non-sparse language model, and test its performance in the few-shot setting.</p>

Author	Title	Link to code	Abstract (short)
Schick and Schütze (2020)	<a href="#">It's Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners</a>	<a href="https://github.com/timoschick/pet">https://github.com/timoschick/pet</a>	<p>We show that performance similar to GPT-3 can be obtained with language models that are much "greener" in that their parameter count is several orders of magnitude smaller. This is achieved by converting textual inputs into cloze questions that contain a task description, combined with gradient-based optimization; exploiting unlabeled data gives further improvements.</p>
Jaegle et al. (2021)	<a href="#">Perceiver IO: A General Architecture for Structured Inputs &amp; Outputs</a>	<a href="https://github.com/deepmind/deepmind-research/tree/master/perceiver">https://github.com/deepmind/deepmind-research/tree/master/perceiver</a>	<p>The recently-proposed Perceiver model obtains good results on several domains (images, audio, multimodal, point clouds) while scaling linearly in compute and memory with the input size. While the Perceiver supports many kinds of inputs, it can only produce very simple outputs such as class scores. <b>Perceiver IO</b> overcomes this limitation without sacrificing the original's appealing properties by learning to flexibly query the model's latent space to produce outputs of arbitrary size and semantics.</p>

## Specific overview

### Speech recognition

Author	Title	Link to code	Abstract (short)
Amodei et al. (2015)	<a href="#">Deep Speech 2: End-to-End Speech Recognition in English and Mandarin</a>	-	We show that an end-to-end deep learning approach can be used to recognize either English or Mandarin Chinese speech—two vastly different languages. Because it replaces entire pipelines of hand-engineered components with neural networks, end-to-end learning allows us to handle a diverse variety of speech including noisy environments, accents and different languages.
Agarwal and Zesch (2019)	<a href="#">German End-to-end Speech Recognition based on DeepSpeech</a>	<a href="https://github.com/AASHISHAG/deepspeech-german">https://github.com/AASHISHAG/deepspeech-german</a>	Description of the process of training German models based on the Mozilla DeepSpeech architecture using publicly available data.

## Information Extraction

Named Entity Recognition

Author	Title	Link to code	Abstract (short)
Anthofer (2017)	<a href="#">A Neural Network for Open Information Extraction from German Text</a>	<a href="https://github.com/danielanthofer/nnoiegt">https://github.com/danielanthofer/nnoiegt</a>	<p>Systems that extract information from natural language texts usually need to consider language-dependent aspects like vocabulary and grammar. Compared to the development of individual systems for different languages, development of multilingual information extraction (IE) systems has the potential to reduce cost and effort. One path towards IE from different languages is to port an IE system from one language to another. PropsDE is an open IE (OIE) system that has been ported from the English system PropS to the German language.</p>
Riedl and Padó (2018)	<a href="#">A Named Entity Recognition Shootout for German</a>	<a href="https://www.ims.uni-stuttgart.de/forschung/ressourcen/werkzeuge/german-ner/">https://www.ims.uni-stuttgart.de/forschung/ressourcen/werkzeuge/german-ner/</a>	<p>We ask how to practically build a model for German named entity recognition (NER) that performs at the state of the art for both contemporary and historical texts, i.e., a big-data and a small-data scenario.</p>
Torge et al. (2021)	<a href="#">Transfer Learning for Domain-Specific Named Entity Recognition in German</a>		<p>Investigation of different transfer learning approaches to recognize unknown domain-specific entities, including the influence on varying training data size.</p>

From: <https://wiki.eolab.de/> - **HSRW EOLab Wiki**

Permanent link: <https://wiki.eolab.de/doku.php?id=ai:nlp:start&rev=1631700306>

Last update: **2021/09/15 12:05**

