Natural Language Processing and Representation Learning, Spring 2020

Logistics:

- **Lecturer**: Marija Stanojevic, SERC 334, marija.stanojevic@temple.edu
- **Class date and time**: ?
- **Location**: Main Campus, ?
- **Office hours**: Tuesday, 12-2pm, SERC 334

Overview: Natural Language Processing (NLP) goal is to create algorithms that will be able to process, understand and communicate using natural languages. It is increasingly popular field due to huge amounts of textual data we are able to collect through the internet and social networks and humongous amount of various knowledge that can be derived from that text. The development of deep learning has improved performance on many natural language problems, putting NLP into spotlight in both academia and industry.

This course is a broad introduction to natural language processing. More attention will be given to modern algorithms which are based on learning representation with deep neural networks. After the course students will be equipped to understand and solve different NLP problems using appropriate algorithms and off-the-shelf solutions.

Course contains a lot of low-demanding homework assignments through which students can deepen their understanding of models and solve real-world problems with existing NLP software. Additionally, students will work on a project. It should solve a real-world problem using two baseline techniques and state-of-art method. Projects will be done in pairs. Result is a 10 min presentation and a 5 page long report with figures that compare the methods. Beside homework and project, students will do a three half-page long written paper reviews and present one of them in the class (7 min).

Requirements:

- Basic knowledge of Machine Learning
- Basic knowledge of Python
- Basic knowledge of Calculus and Linear Algebra
- Basic knowledge of Probability and Statistics

Books (optional):

- Jurafsky, Dan, and James H. Martin. "Speech and language processing." 3nd edition draft. (2018).
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016. (TU library)
- Manning, Christopher D., Christopher D. Manning, and Hinrich Schütze. *Foundations of statistical natural language processing*. MIT press, 1999. (available online through TU library)

Course schedule:

Week	Topics	Additional Materials (not required to go through)	Homeworks and Projects
1	-Introduction to NLP -Text preprocessing -Rule based models -Stemming / lemmatization -Introduction to Tensorflow	-Ch 1. Foundations of Statistical Natural Language Processing -Ch 1 and 2. Speech and Language Processing -https://www.tensorflow.org/	-Projects out
2	-Naive Bayes -N-grams -Hidden Markov Models -Part-of-Speech Tagging	-Ch 5, 6, 9 and 10. Foundations of Statistical Natural Language Processing -Ch 3, 4 and 8. Speech and Language Processing	-HW 1 out (preprocessing, tfidf, LSI) -Projects selection due
3	-Tf-idf model -Latent Semantic Indexing -Projects proposals review	-Ch 15. Foundations of Statistical Natural Language Processing - Ch 6. Speech and Language Processing	-HW 1 due -HW 2 out (n-grams, HMM, POS)
4	-Paper reviews -Vector Space Model -Word2vec -Glove	-Ch 6, 14. Speech and Language Processing	-HW 2 due -HW 3 out (w2v, d2v, glove)
5	-Paper reviews -Representation of Sentence Meaning -Doc2vec -Recurrent Neural Networks -LSTM, bi-LSTM, GRU	-Ch 10. Deep Learning -Ch 9. Speech and Language Processing	-HW 3 due -HW 4 out (seq-2-seq)
6	-Paper reviews -Sequence-to-Sequence Architectures -Attention Modeling -Self-attention, Transformers	-Ch 9. Deep Learning -Attention is all you need	-Project, part 1 due
7	-Paper reviews -Syntactic Parsing -Statistical Parsing -Dependency Parsing -Constituency Parsing -Semantic Parsing	-Ch 11 and 12. Foundations of Statistical Natural Language -Ch 10-13. Speech and Language Processing	-HW 4 due -HW 5 out (parsing)

8	-Paper reviews -Language modeling -Pre-training -BERT -Self-supervised learning -ALBERT	- BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding - ALBERT: A Lite BERT for Self-supervised Learning of Language Representations	-HW 5 due -HW 6 out (BERT)
9	-Paper reviews -Information Extraction -Semantic Role Labeling -Text Classification -Sentiment Analysis	-Ch 17-19. Speech and Language Processing	-HW 6 due -HW 7 out (text classification)
10	-Paper reviews -Word sense disambiguation -Coreference Resolution and Entity Linking -Discourse Coherence -Question Answering	-Ch 7. Foundations of Statistical Natural Language -Ch 20, 21, 23. Speech and Language Processing	-HW 7 due -HW 8 out (question answering) - Project part 2 due
11	-Paper reviews -Text Generation -Dialog Systems and Chatbots	-Ch. 24, 25. Speech and Language Processing	-HW 8 due -HW 9 out (chatbot)
12	-Paper reviews -Statistical Machine Translation -Neural Machine Translation	-Ch 13. Foundations of Statistical Natural Language -Ch 22. Speech and Language Processing	-HW 9 due -HW 10 out (NMT)
13	-Project presentations		
14	-Project presentations		-HW 10 due - Project part 3 due

Grading:

- 40% Homeworks (10 homework assignments, 4% each)
- 10% Written papers reviews (2 reviews, 5% each): to receive 5% points, student must show criticism; to receive 4% points student must show understanding of methodology.
- 10% Oral review (7 min long presentation in the class)
- 40% Project (part 1 - 5%, part 2 - 7%, part 3 - 10%, presentation - 5%; report - 13%)