Tianyu General Exam Reading List

Textbooks

1. Speech and Language Processing (3rd edition), by Daniel Jurafsky and James H. Martin.
    a. Ch. 3: N-gram Language Models.
    b. Ch. 7: Neural Networks and Neural Language Models.
    c. Ch. 9: Deep Learning Architectures for Sequence Processing.
    d. Ch. 11: Transfer Learning with Contextual Embeddings and Pre-trained language models.

Papers

Pre-trained language models

1. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. Advances in Neural Information Processing Systems (NIPS), 30.
2. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), pages 41714186.
3. Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tieyan Liu. 2019. **Representation degeneration problem in training natural language generation models**. In International Conference on Learning Representations (ICLR).

Sentence embeddings

4. Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2015. **Skip-thought vectors**. In Advances in Neural Information Processing Systems (NIPS), pages 3294–3302.
5. Lajanugen Logeswaran and Honglak Lee. 2018. **An efficient framework for learning sentence representations**. In International Conference on Learning Representations (ICLR).
6. Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. **A simple but tough-to-beat baseline for sentence embeddings**. In International Conference on Learning Representations (ICLR).

7. Nils Reimers and Iryna Gurevych. 2019. **SentenceBERT: Sentence embeddings using Siamese BERTnetworks**. In Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLPIJCNLP), pages 3982–3992.
8. Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. **On the sentence embeddings from pre-trained language models**. In Empirical Methods in Natural Language Processing (EMNLP), pages 9119–9130.

Contrastive learning

9. Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. 2014. **Discriminative unsupervised feature learning with convolutional neural networks**. In Advances in Neural Information Processing Systems (NIPS), volume 27.
10. Saunshi, Nikunj, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. 2019. **A theoretical analysis of contrastive unsupervised representation learning**. In International Conference on Machine Learning (ICML), pages 5628-5637.
11. Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. **A simple framework for contrastive learning of visual representations**. In International Conference on Machine Learning (ICML), pages 1597–1607.