

Long-term Objective

To contribute towards solving Natural Language Understanding (NLU), which is a crucial aspect of AI. To conduct research on machine/deep learning models and algorithms which are able to discover, capture, represent, generalize and transfer (meta)understanding for and across tasks of human cognition involving NLU.

Natural Language research has traditionally focused on specialized tasks for modelling and extracting linguistic regularities from textual data, more specifically, it has focused on pipelined *text* processing systems involving multiple individual components each responsible for one type of linguistic representation. The recent resurgence of deep learning has not only completely revamped the traditional linguistic processing components by replacing manual feature engineering (where neural networks in various architectural configurations are used as feature extractors in standard natural language text processing tasks), it has also been applied to create more robust end-to-end natural language *understanding* systems.

Due to its flexibility and adaptability, deep learning opens new avenues for NLU research that potentially can lead to breakthroughs in the following three interrelated areas that align with my research interests.

First, humans do not learn languages solely by reading *static* text – it is a well-known fact that one can learn to speak a language without understanding it in its written form – an ability also possessed by the illiterate and blind. Instead, language learning is *dynamic*, situated and connected with vision, sound, actions, and interactions with the environment. How do we design neural network models and associated learning algorithms to approximate these underlying learning mechanisms and use them to discover, capture, represent, generalize and transfer the semantic knowledge across different tasks? Imagining a neural network model that is able to watch a movie as well as reading a book telling the same story as the movie, it should be capable of generating a unified semantic understanding of the story, and apply this understanding to other tasks that do or do not involve explicit natural language text or utterances. This research potentially connects reinforcement, multi-task, as well as unsupervised learning, and can involve inventing new variants of neural networks that have novel memory and reasoning components (Santoro et al., 2017; Parisotto and Salakhutdinov, 2018).

Second, because the amount of textual information available has been ever increasing, I am also interested in deep learning models for linguistic structured prediction (which is the primary theme of natural language text processing), and combining it with grounded language learning. One promising line of research is neural network models with built-in mechanisms to deal with inductive biases and structured prediction for text processing. It has been shown repeatedly that neural network models incorporating structured learning mechanisms outperform their vanilla counterparts, and combining structured learning with deep learning is an emerging theme, especially in devising both empirical and formal methods that faithfully take into account the respective neural models, instead of relying upon techniques originally developed for other models (Watanabe and Sumita, 2015; Weiss et al., 2015; Andor et al., 2016; Lample et al., 2016; Lee et al., 2016). I barely scratched the surface of investigating this topic in my Ph.D. thesis for recurrent neural networks, and hope to further this not just by adding external constraints to existing neural networks, but by designing neural networks that are able to handle inductive biases and structured prediction naturally, as part of their internal architecture (Belanger et al., 2017; Sabour et al., 2017; Hinton et al., 2018; Tu and Gimpel, 2018).

Finally, there remains a lot of work to be done in coming up with novel architectures and learning algorithms for language processing tasks of high practical relevance, such as machine translation (He et al., 2016; Vaswani et al., 2017; Xia et al., 2017; Zhang et al., 2018). I also find this type of work exciting, and envision a future where performance of these systems can be fundamentally improved by the results of all aforementioned lines of research.

March. 2018

References

- Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. Globally normalized transition-based neural networks. In *Proc. of ACL*, 2016.
- David Belanger, Bishan Yang, and Andrew McCallum. End-to-end learning for structured prediction energy networks. In *Proc. of ICML*, 2017.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tieyan Liu, and Wei-Ying Ma. Dual learning for machine translation. In *Proc. of NIPS*, 2016.
- Geoffrey E. Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with EM routing. In *Proc. of ICLR*, 2018.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. Neural architectures for named entity recognition. In *Proc. of NAACL*, 2016.
- Kenton Lee, Mike Lewis, and Luke Zettlemoyer. Global neural CCG parsing with optimality guarantees. In *Proc. of EMNLP*, 2016.
- Emilio Parisotto and Ruslan Salakhutdinov. Neural map: Structured memory for deep reinforcement learning. In *ICLR*, 2018.
- Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. Dynamic routing between capsules. In *Proc. of NIPS*, 2017.
- Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Tim Lillicrap. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, 2017.
- Lifu Tu and Kevin Gimpel. Learning approximate inference networks for structured prediction. In *Proc. of ICLR*, 2018.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proc. of NIPS*, 2017.
- Taro Watanabe and Eiichiro Sumita. Transition-based neural constituent parsing. In *Proc. of ACL*, 2015.

- David Weiss, Chris Alberti, Michael Collins, and Slav Petrov. Structured training for neural network transition-based parsing. In *Proc. of ACL*, 2015.
- Yingce Xia, Fei Tian, Lijun Wu, Jianxin Lin, Tao Qin, Nenghai Yu, and Tie-Yan Liu. Deliberation networks: Sequence generation beyond one-pass decoding. In *Proc. of NIPS*, 2017.
- Zhirui Zhang, Shujie Liu, Mu Li, Ming Zhou, and Enhong Chen. Joint training for neural machine translation models with monolingual data. *arXiv preprint arXiv:1803.00353*, 2018.