



Learning to Generate High-quality Dialogue

Motivation

Applications of open-domain conversational agents are becoming widespread. However, training such agents to generate high-quality responses is still a big challenge as the quality of responses depends on various factors. Recent methods train agents directly by gold responses from training sets. These methods have been shown generating low-quality responses at evaluation. In this thesis, we propose to train a function that quantifies the quality of the generated responses by a deep preference learning method. Then, we use this function as a reward estimator in a reinforcement learning model to train agents.

Task Description

- Preparing a few existing agents for comparison
- Training a reward estimator
- Training an RL-based agent using the reward estimator
- Comparing with the prepared agents



References

- See Abigail, et al., (2019). What makes a good conversation? How controllable attributes affect human judgments. In Proc. of NAACL-HLT 2019, pp. 1702-1723.
- Li Jiwei et al., (2016). Deep reinforcement learning for dialogue generation. In Proc. of EMNLP 2016. pp. 1192-1202.

Contact

Analysis



Programming



Literature



Prof. Dr. Iryna Gurevych
Mohsen Mesgar
Dr. Edwin Simpson

thesis@ukp.informatik.tu-darmstadt.de