Nicebot: The Sentiment Analysis Chatbot

Yoni Lerner
yonlern@stanford.edu

Problem

Many companies are using chatbots in commercial applications today, and users can often become frustrated with such bots. Some companies, such as Dashbot[1], already offer sentiment analysis as a diagnostic tool for chatbot success. Nicebot takes a different approach, directly training to maximize user happiness, as measured by the sentiment of the user's response.

Datasets and Features

The sentiment analysis model was trained using the IMDB movie review dataset.[2] The Seq2Seq model along with the discriminator was trained using the Cornell movie dataset, which contains conversations from movie scripts.[3] The movie set is unlabeled, but labeled triplets are extracted from the conversations. The model and discriminator both accept a sequence of word embeddings, which allow the model to train faster and better generalize to uncommon words.

Results

Though no concrete model was trained, the contribution of this project is in its framework. The architecture used is able to learn to produce positive conversational responses despite using one of which is unlabeled for positivity and one of which is non-conversational. This paradigm could be used in other applications were conventional supervised datasets are unavailable.

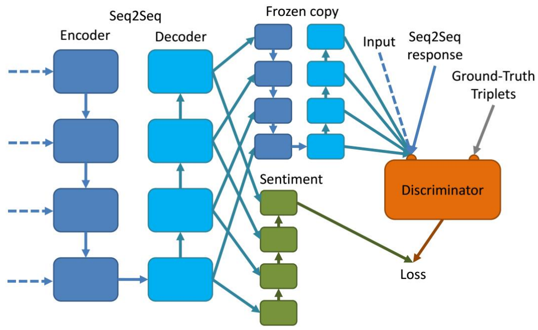

Model Architecture

The model is based around a standard Seq2Seq query/response network. However, the model is not trained directly on query/response pairs. Instead, when the model outputs a response, the response is fed to a duplicate frozen model as an input, and a follow-up response is generated, creating a query/response/follow-up triplet. This triplet is given to a GAN-style discriminator, which attempts to discriminate it from real triplets sampled at random from a dataset of real two-actor conversations. The follow-up is propagated through a frozen pre-trained sentiment analysis network. The model is then trained to both maximize positive sentiment of the follow-up and fool the discriminator (using standard GAN techniques). Since the model is essentially trained using self-play, the GAN is necessary to prevent it from ”hacking” the sentiment by only producing positive responses.

Discussion

The problem as stated turned out to be significantly more difficult than expected. The model went through many iterations, beginning as a reinforcement-learning algorithm that used sentiment analysis to generate a reward. However, it was realized that this was essentially performing inverse reinforcement learning, and that if the sentiment analyzer is taken as ground-truth, it is more effective to simply backpropagate through it rather than use reinforcement learning. In the end, this project was far too complex and computationally demanding for a one-paper team to complete in ten weeks, but much was learned along the way, and the core idea seems to still be valid.

Future

The main challenge in this project was the amount of time and computation required for implementing and training such a complex model. In the future, the full architecture would be trained and refined to evaluate its potential.

References

[1]"Dashbot Conversational Analytics - Dashbot", *Dashbot*, 2018. [Online]. Available: https://www.dashbot.io/. [Accessed: 20- Mar- 2018].
[2]A. Maas, R. Daly, P. Pham, D. Huang, A. Ng and C. Potts, "Learning Word Vectors for Sentiment Analysis", *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142-150, 2011.
[3]C. Danescu-Niculescu-Mizil and L. Lee, "Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs.", *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics, ACL 2011*, 2011.