Abstract

Development of trainable policies for spoken dialogue systems

Thanh Le

In human-human interaction, speech is the most natural and effective manner of communication. Spoken Dialogue Systems (SDS) have been trying to bring that high level interaction to computer systems, so with SDS, you could talk to machines rather than learn to use mouse and keyboard for performing a task. However, as inaccuracy in speech recognition and inherent ambiguity in spoken language, the dialogue state (user's desire) can never be known with certainty, and therefore, building such a SDS is not trivial.

Statistical approaches have been proposed to deal with these uncertainties by maintaining a probability distribution over every possible dialogue state. Based on these distributions, the system learns how to interact with users, somehow to achieve the final goal in the most effective manner. In Reinforcement Learning (RL), the learning process is understood as optimizing a policy of choosing action conditioned on the current belief state.

Since the space of dialogue states is not small, even with a very little SDS, a typical RL algorithm must experience hundreds of thousands dialogues to find the optimal policy. Recently, the GP-Sarsa algorithm has leveraged the learning process with the power of Gaussian Processes (GP). One has already shown that the GP-Sarsa algorithm might requires only thousands of dialogues to get the optimal policy.

This thesis further examines the GP-Sarsa algorithm in SDS, reimplementing and evaluating it in a real world scale problem. As a well-known requirement of gathering experience with RL, a domain-independent user simulator for SDS is also investigated in the thesis.

Keywords: POMDP, Bayesian methods, HMM, dialogue systems, NLP