

Analyzing Language Learned by an Active Question Answering Agent

We analyze the language learned by an agent trained with reinforcement learning as a component of the ActiveQA system [Buck et al., 2017]. In ActiveQA, question answering is framed as a reinforcement learning task in which an agent sits between the user and a black box question-answering system. The agent learns to reformulate the user's questions to elicit the optimal answers. It probes the system with many versions of a question that are generated via a sequence-to-sequence question reformulation model, then aggregates the returned evidence to find the best answer.

Author: CHRISTIAN , Buck (Google)

Presenter: CHRISTIAN , Buck (Google)