Research Statement

He He
Stanford University

June 16, 2017

The future of virtual assistants, self-driving cars, and smart homes requires intelligent agents to work in real-time, dynamic environments and to serve users in an intimately, communicative way. Language plays a key role in these innovations as a channel to interact with humans. Traditional natural language processing tasks and methods have focused on static or batch settings where all inputs are ready and the model makes a one-time prediction, e.g., text classification, parsing, and machine translation, However, in a dynamic environment, data (inputs) can arrive sequentially and the model needs to act on those incomplete data in an online fashion. Furthermore, when a human is involved, the model should seek information from users when necessary, for example, asking for their preference, clarifying intentions, or coordinating actions.

My goal is to build models that can interact with its environment through decision-making and (natural language) communication, especially by applying techniques from reinforcement learning and imitation learning. On the machine learning side, I am interested in questions such as: How can a model make predictions given incomplete information, and acquire costly information according to its needs? How can we solve complex, combinatorial problems in a sequential way by trading off accuracy for speed? On the application side, I have worked on task-oriented dialogue systems where the agent (model) needs to communicate and collaborate with its human partners, as well as question answering games where the model competes against human players.

Sequential Question Answering. Quiz bowl is a trivia game where players compete to answer a paragraph-long question before the entire question is revealed. These questions are written so that they can be interrupted by someone who knows more about the answer; that is, harder clues are at the start of the question and easier clues are at the end of the question. It is related to core challenges in natural language processing as mentioned above: The model needs to make predictions with only a partial question; it also needs to decide when to give the answer by estimating the opponent's strength and evaluating a tradeoff between speed and accuracy. I applied imitation learning algorithms [1] to decide where to answer in the middle of a question, and later, a deep reinforcement learning network [2] that models different opponents and plays adaptively. Our quiz bowl system received the best demonstration award at NIPS 2015.

Symmetric Dialogue. Language is the most natural way of communication for humans, and a dialogue interface empowers models to exchange information with users. Most current task-oriented dialogue systems focus on virtual assistance settings such as restaurant searching and flight booking, which consists of mainly information lookup and query answering. We proposed a more symmetric dialogue setting [3], where the user and the dialogue agent must communicate to achieve a common goal given private knowledge bases. We also developed a novel neural network model that bridges two main approaches to dialogue: the fully structured semantic parsing approach and the unstructured end-to-end neural generation approach. Our model successfully generated natural, human-like utterances, and learned to exchange information strategically with humans. This work aims to build collaborative agents in a relatively structured context (i.e. knowledge bases). Moving towards more realistic and nuanced settings, we are currently working on negotiation, where decision-making is rendered by much more complex and interesting language, e.g. persuading, deceiving, asserting.

In future, I am interested in exploring how can the model continue to learn during interaction with humans and how can multiple agents learn from each other in a shared environment. In addition, data collection and evaluation are still not well-defined for interaction tasks, which hinders progress in this area. I hope to work on building better crowdsourcing platforms for efficient data collection and human evaluation.

References

[1] Jordan Boyd-Graber, Brianna Satinoff, He He, and Hal Daum'e III. Besting the quiz master: Crowdsourcing incremental classification games. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2012.

[2] He He, Jordan Boyd-Graber, Kevin Kwok, and Hal Daum'e III. Opponent modeling in deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2016.

[3] He He, Anusha Balakrishnan, Mihail Eric, and Percy Liang. Learning symmetric collaborative dialogue agents with dynamic knowledge graph embeddings. In *Association for Computational Linguistics (ACL)*, 2017.