

Task-oriented dialog (TOD) systems converse with users to help them with specific tasks such as calendar enquiry, restaurant reservations, and tourist package recommendations. To achieve task success, TOD systems must interact with a task-specific knowledge source. For example, a restaurant reservation system has to be grounded on a knowledge base that contains the names of restaurants, and their details. End-to-end TOD systems are a type of TOD system that is trained using past human-to-human dialogs and do not require any handcrafted state representations and their corresponding annotations on each dialog.

Several challenges prevent end-to-end TOD systems from being used for real-world applications. In this thesis, we address three such challenges. The common theme across the contributions is that the proposed TOD systems that are aimed at tackling these challenges generalize well to knowledge unseen during train. Firstly, to improve the performance of end-to-end TOD systems grounded on structured knowledge bases, we propose three novel neural models -- BoSsNet, MLM-Net and CDNet, (2) To reduce the training cost, we define and study the novel problems of training end-to-end TOD systems without the need for explicit KB query annotations, and (3) to extend the types of knowledge sources supported by end-to-end TOD systems, we define and study the novel task of end-to-end learning of a TOD system that troubleshoots user's problems by using a flowchart and a corpus of FAQs.