

Travellers often post questions online to seek personalized travel recommendations by describing their preferences and constraints with respect to locations, points of interests, budget, etc. They also, at times, post queries asking for comparisons between cities, tourism sites, etc when making their travel plans. In this thesis we study the novel tasks of answering such recommendation and comparison questions from the tourism domain.

We focus our attention on a class of recommendation questions that seek entities. We refer to them as Multi-sentence entity-seeking recommendation questions (MSRQs) i.e., questions that expect one or more entities as an answer. In the tourism domain, such entities can occur in the form of Points-of-Interest (POIs); e.g, names of hotels, restaurants, tourist sites. We answer entity-seeking recommendation questions in two settings: (i) QA with intermediate annotations (ii) QA without intermediate annotations. In each setting we formulate a new problem and create new datasets which we hope will help further research in QA. In the first setting, we develop a pipelined model which breaks down the task of question-answering into a question-parsing task followed by knowledge-base querying. Learning a question-parser requires large amounts of training data and we overcome this challenge by employing a constraint driven learning framework that uses a small set of expert-annotated questions, along with a larger set of crowd-sourced partially-annotated questions.

In contrast to the first setting, in the second approach we use a collection of reviews to directly answer questions, without explicitly parsing questions. Answering such questions poses novel challenges of reasoning at scale, since review collections for each entity can be very large, noisy, contain subjective opinions, and each question can have thousands of entities to choose from to return as ‘possible answers’. Further, existing retriever-ranker strategies commonly employed to address challenges of scale do not work on this task. In response, we present a cluster-retrieve-rerank architecture that helps address some of these challenges. Additionally, in order to accommodate reasoning over physical locations of entities, we extend this work by developing a joint spatio-textual model. We develop a modular spatial-reasoning network that uses geo-coordinates of location names mentioned in a question, and, of candidate answer entities, to reason over only spatial constraints. We combine the spatial-reasoner with the textual QA system to develop a joint spatio-textual QA model. We demonstrate that our joint spatio-textual model performs significantly better than models employing only spatial or textual reasoning.

Lastly, we also study the problem of answering comparison questions. We define a novel task of generating entity comparisons from textual corpora in which each document describes one entity at a time. We generate entity comparisons in a tabular form in which attribute-value phrases, opinion phrases, and other descriptions are clustered and organized topically, thus, allowing for direct comparisons. Our tabular summaries balance information about the entities being compared and in our user studies we find that balanced clusters are strongly preferred and users acquire as much information about the entities, by using the tables, as they do using articles.