Researchers develop a system that can recommend personalized and healthy recipes

January 21 2021, by Ingrid Fadelli

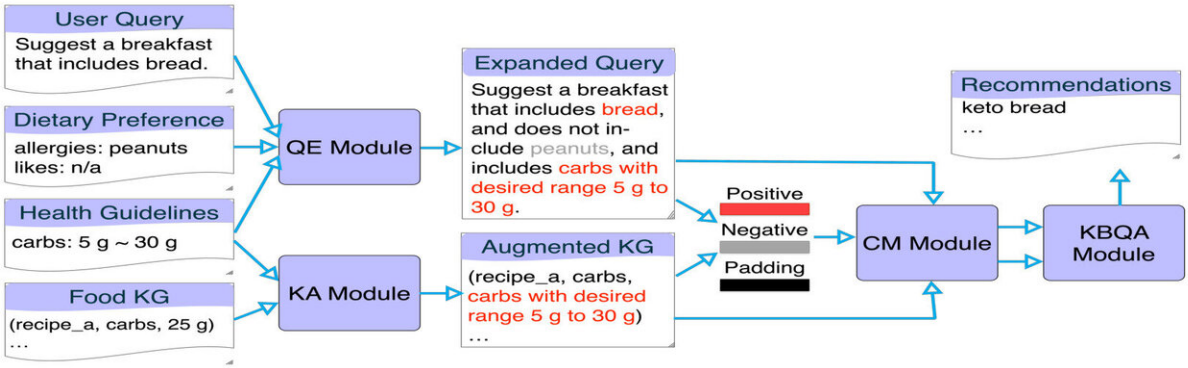

Figure summarising the researchers' personalized food recommendation architecture: Given a user query in natural language, along with the user's context (comprising dietary preferences and health guidelines), the system finds top-ranked and personalized answers from the FoodKG food knowledge graph. QE, KA and CM stand for Query Expansion, Knowledge Graph Augmentation and Constraint Modeling, respectively. Positive constraints are in red, negative constraints in gray, and non-constraints in black. KBQA refers to the underlying knowledge graph question answering system. Credit: Chen et al.

Researchers at Rensselaer Polytechnic Institute and IBM Research in New York have recently created pFoodReQ, a system that can recommend recipes tailored around the preferences and dietary needs of individual users. This system is outlined in a paper pre-published on arXiv and set to be presented at the 14th International Conference on

Web Search and Data Mining (WSDM) in March.

"Our work focuses on personalized food recommendation," Mohammed J. Zaki, one of the researchers who developed the system, told TechXplore. "In particular, given a user query in natural language, we want to retrieve the top matches in a recipe dataset."

The short-term goal of the study carried out by Zaki and his colleagues was to help people find healthy recipes that satisfy both their dietary needs and inclinations. Ideally, they wanted these recipes to reflect both a users' intent (i.e., whether they are planning a quick or more elaborate meal), their preferences (i.e., whether they like or dislike specific ingredients or cuisines), dietary constraints (i.e., if they are allergic or intolerant to certain products) and health status (i.e., if they are diabetic, overweight, or have specific health conditions).

"The key idea is that given the same query, the response should actually be different for different users," Zaki said. "In other words, the responses should be personalized. This is a very challenging task, especially in terms of determining the implicit constraints that are actually relevant to the query."

A broader objective of the research by Zaki and his colleagues was to combine machine learning methods with semantics, the branch of linguistics related to meaning. To do this, the researchers combined deep-learning techniques for answering questions with FoodKG, a large-scale food knowledge graph they developed. The FoodKG graph, which is publicly available online, contains over 67 million records (or triples). These records include approximately one million recipes, graph representations of the relationships between these recipes and the ingredients required to complete them, as well as data related to the properties of ingredients, nutrient contents and different cuisines or food categories.

"Our framework has a number of unique advantages compared to existing works focusing on food recommendation," Zaki explained. "First, our aim is to cast the personalized recipe recommendation task as question answering in natural language, for ease of use. Second, our system can handle explicit requirements that mention food tags (e.g., Italian dishes) and allowed ingredients (e.g., must contain spinach), as well as negations (e.g., should not contain walnuts)."

Negations, or statements that highlight things that should not be included in a recommended item, are very difficult for deep-learning systems to tackle. This is mainly because deep neural networks are not always good at distinguishing between positive and negative statements (e.g., a recipe must contain peanuts vs. it must not contain peanuts). The researchers were able to overcome this limitation by introducing specific positive/negative cues for the neural network.

Another unique characteristic of the food recommendation system is that it captures both a user's unique qualities or preferences and general guidelines for healthy nutrition. This allows it to provide personalized recipe recommendations that also promote good nutrition and health.

"To achieve this, we leveraged a semantics-based approach by modeling the user preferences and health guidelines (e.g., the total fat consumption for all adults should be 20-35% of total calorie intake, from the American Diabetes Association lifestyle guidelines) as knowledge subgraphs that are combined with the relevant FoodKG subgraph to retrieve good suggestions."

The framework is adaptive and can learn from a user's past logs. This means that it acquires information about a person's dietary preferences and needs over time, then uses this information to suggest new, healthy recipes aligned with these preferences and needs.

"The under-the-hood technical approach behind our recommendation system relies on three modules that leverage the underlying question-answering deep-learning system," Zaki said. "The first and most important one is the query expansion module, where the original natural language query is expanded to include the implicit constraints from the user and health knowledge subgraphs so that the final query sent to the question-answering module has all requirements and constraints made explicit."

The framework's second component, the constraint modeling module, handles both positive and negative ingredient constraints (e.g., must contain or must not contain peanuts). The last module, on the other hand, performs what is known as knowledge graph augmentation.

"Deep learning systems struggle with numeric comparison constraints (e.g., dishes that are under 500 calories)," Zaki explained. "We solve this by converting the numeric comparison constraints by dynamically 're-writing' or augmenting the FoodKG subgraph into Boolean constraints (true/false) corresponding to the given numeric constraints (e.g., recipes that meet the constraint are dynamically marked as true)."

Combined, the framework's three modules allow it to give healthy and personalized recipe recommendations. In addition to this highly performing food recommendation system, the researchers created a benchmark of natural language questions and ground truth answers related to food suggestions. Remarkably, before this study, such a benchmark did not exist. Moreover, the questions contained in it are styled after real food-related questions posted by users on Reddit.

To evaluate their framework and compare it with existing food recommendation systems, the researchers carried out a user study in which a few participants were asked to evaluate the suggestions of different systems, including their own. All the systems based their

recommendations on the participants' responses to 50 questions, which explored their food-related preferences and dietary needs. Remarkably, the framework created by Zaki and his colleagues outperformed all the techniques to which it was compared, with most participants expressing their preference for the recommendations it had provided over those produced by other systems.

"We hope that the public availability of this benchmark will spur much-needed follow-on work and future improvements," Zaki said. "On the technical side, we show that our proposed approach of combining query expansion, constraint modeling and knowledge graph augmentation with our state of the art deep-learning based knowledge graph question answering system (BAMnet) provides tremendous improvement in food recommendation compared to existing approaches."

The research carried out by Zaki and his colleagues was funded by IBM Research's AI Horizons Network and is part of a broader project called RPI-IBM HEALS, which stands for Health Empowerment by Analytics, Learning, and Semantics. The code of the food recommendation system they developed is available online at: https://github.com/hugochan/PFoodReq.

In the future, this paper could inform the development of digital assistants that help users to make healthier food choices and source new recipes. Zaki and his colleagues are now working on improving their system in several ways.

"The existing system should be considered as a first step in personalized food recommendation, but many challenges still remain," Zaki said. "We need even more complex question-answering benchmarks that deal with the implicit user intents and diverse scenarios. Recommending healthier alternatives and potential substitutions also needs more work. Handling more complex constraints and finding the most relevant set of

constraints that should be considered when answering a query remains a challenge and this can also be considered as incorporating common-sense reasoning into the framework."

More information: Personalized food recommendation as constrained question answering over a large-scale food knowledge graph. arXiv:2101.01775 [cs.CL]. arxiv.org/abs/2101.01775