Abstract

Overweight and obesity affect billions of adults. Limiting the scope to nutrition, help must offer diet guidance that fits easily into daily life. Health-aware meal plan recommendations must integrate user preferences, nutrition goals, and everyday context, like time and date. However, existing recipe data sets and recommender systems lack general health guidelines, they miss capturing user-recipe interactions, or do not provide explanations or other persuasion methods to improve on the recommendations effectiveness. This thesis closes these gaps and evaluates
proposed systems for health-aware meal planning. We give an extensive background on nutrition and persuasion, an overview of LLMs, and introduce related work in terms of data sets and recommendation approaches. Then, our contribution can be split into three parts.

Our first contribution is *HUMMUS*, a large, linked food graph with approximately 500 000 recipes, 300 000 users, 1.9 million interactions, nutrition scores, and semantic links to *FoodOn* (a food product hierarchy) and *FoodData* (nutritional ingredient and food product data). It enables reasoning over ingredient classes and nutrient checks. Its size, presence of interactions, and diverse nutrition information improve usability when compared to related data sets. We discuss sparsity, preprocessing, and filtering options that support different experiments.

Our second contribution is a study of food-centric behaviour. We collected questionnaire data and meal logs with an online app to explore correlations between *food skills (FS)*, *cooking skills (CS)*, and intake context. The sample includes 78 participants. Results show that FS and CS correlate. These findings and the app help to explore the domain of food behaviour and to improve further studies and recommender systems.

Our third contribution is a framework for a daily meal plan recommendation, including explanations. We design and implement two approaches that accept natural language input. Our first approach, a *KBQA* system, recommends daily meal plans, supports multiple nutrient constraints, and produces path-grounded explanations. It expands its baseline *PFoodReq* by those features, but also inherits scalability issues and enforces soft constraints. We then propose the *FoodRAG* architecture, an LLM-based, agentic, modular pipeline that performs recipe retrieval, validation, ingredient substitution, and step-wise explanation. Despite a higher computational cost than the baselines, its extensive adaptability produces more targeted meal-plan recommendations. Both systems are evaluated and compared to baseline recommendation approaches (collaborative and content-based). Persuasion is addressed by integrating fitting gamification methods, providing explanation generation as well as generic UI aids. As collaborative methods underperform on sparse data like *HUMMUS*, we evaluate simple content-based baselines, the *KBQA* system, and three *FoodRAG* variants. The *Text-to-Pandas* retriever (an LLM with natural language input questions to generate *Pandas* queries used on the *FoodRAG* data set) gives the strongest recipe retrieval results.

Keywords: Health-aware recommendation, Meal planning, Knowledge graphs, Recipe data set, Explainability, Persuasion, KBQA, RAG.