

Knowledge-Powered Recommendation for an Improved Diet Water Footprint

Saurav Joshi¹, Filip Ilievski^{1,2}, Jay Pujara¹

¹USC Information Sciences Institute, Marina del Rey, CA, USA

²Vrije Universiteit, Amsterdam, The Netherlands
syjoshi@isi.edu, f.ilievski@vu.nl, jpujara@isi.edu

Abstract

According to WWF, 1.1 billion people lack access to water, and 2.7 billion experience water scarcity at least one month a year. By 2025, two-thirds of the world’s population may be facing water shortages. This highlights the urgency of managing water usage efficiently, especially in water-intensive sectors like food. This paper proposes a recommendation engine, powered by knowledge graphs, aiming to facilitate sustainable and healthy food consumption. The engine recommends ingredient substitutes in user recipes that improve nutritional value and reduce environmental impact, particularly water footprint. The system architecture includes source identification, information extraction, schema alignment, knowledge graph construction, and user interface development. The research offers a promising tool for promoting healthier eating habits and contributing to water conservation efforts.

Introduction

Water scarcity is a prevailing issue that poses significant environmental concerns (Rijsberman 2006). This essential life-giving resource is critical not just for humans but also for the planet at large. Consequently, comprehending and regulating its consumption has become a crucial responsibility in the contemporary context. A key initiative in this realm involves the development of the metric known as the *water footprint* which measures and raises consciousness about water usage (Hoekstra et al. 2009).

The global food sector is a significant contributor to climate change (Vermeulen, Campbell, and Ingram 2012), biodiversity loss and land-use change (Foley et al. 2005), and rampant exhaustion of freshwater resources (Wada et al. 2010). A comprehensive study (Mekonnen and Gerbens-Leenes 2020) indicates that the food industry is one of the most water-intensive sectors. To effectively manage water usage, numerous studies have been directed to establish dietary guidelines that advocate a more sustainable approach to food consumption with minimal environmental impacts (Springmann et al. 2018; Mekonnen and Fulton 2018). However, despite the clear importance of healthy eating, many struggle with it due to busy lifestyles or lack of meal planning motivation (van Pinxteren, Geleijnse, and Kamsteeg 2011). Health-aware food recommender systems use both

content-based approaches, derived from user and item attributes (Harvey, Ludwig, and Elsweiler 2013), and collaborative filtering, which bases predictions on historical community preferences (Ge et al. 2015). In (Liu et al. 2019), authors explore knowledge graphs as a solution for recommendation systems. Moreover, in recipe recommendations, these graphs view recipes as feature aggregates, enhancing prediction accuracy (da Silva 2021).

Leveraging these advancements, this paper proposes the development of a recommendation engine powered by knowledge graphs. It aims to help users navigate their diet planning efforts more sustainably by proposing substitutes for food ingredients in user’s recipes that offer improved nutritional value and reduce the environmental impact regarding sustainability and water footprint. This concept stems from the ability of the knowledge graphs to offer context-aware recommendations, reveal hidden or implicit relationships, and facilitate informed decision-making; thereby serving as an effective tool for users to eat healthily and sustainably.

System Architecture

As seen in Figure 1, the system architecture comprises five sequential steps: Source Identification, Information Extraction, Schema Alignment, Knowledge Graph Construction, and User Interface Development.

Source Identification. Our research employs the FoodKG resource (Haussmann et al. 2019), a comprehensive database spanning over a million recipes, ingredients, and nutrients, encompassing approximately 67 million data points. We interpret and process information from three primary sources to lay the foundation for FoodKG: the Recipe1M dataset (Marin et al. 2021), the USDA¹ database which catalogs the nutritional content of each ingredient, and FOODON which is an extensive food ontology system for ingredient organization (Dooley et al. 2018). Additionally, we incorporate complementary data on water footprints from academic journals and papers publicly available as PDF documents.

Information Extraction. Our approach involves the application of the Knowledge Graph Toolkit (KGTK) (Ilievski et al. 2021). We transform the information from the

¹<https://www.usda.gov/topics/food-and-nutrition>

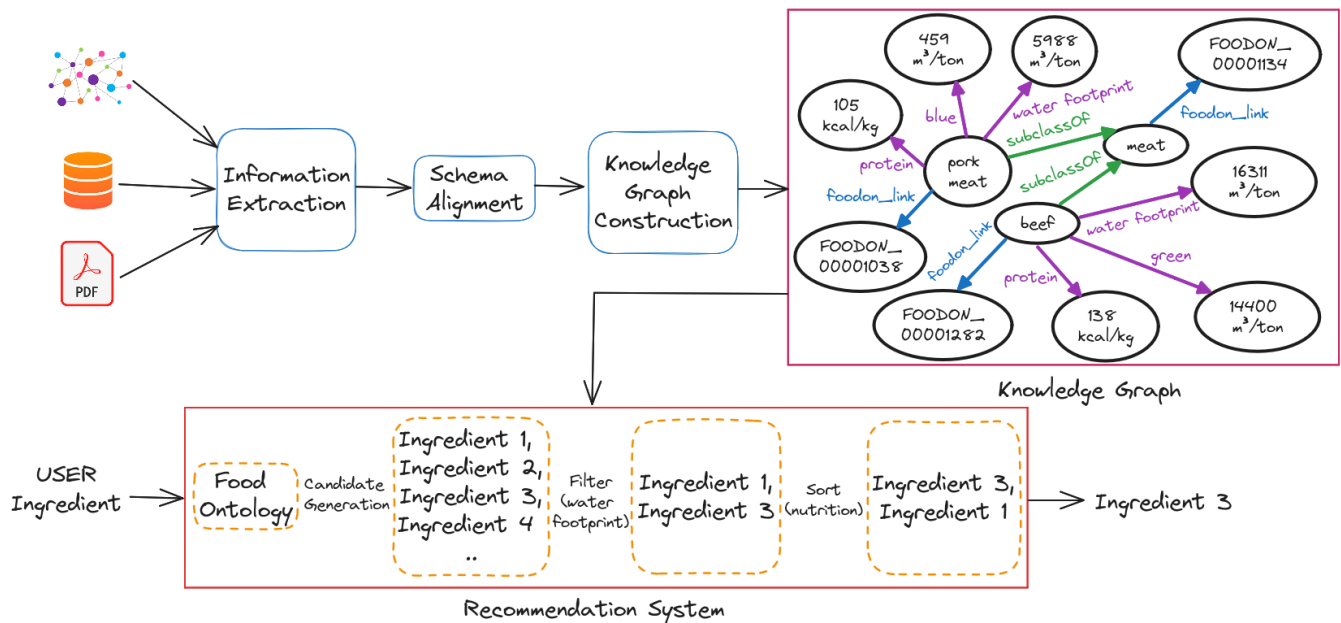


Figure 1: System Architecture consisting of five steps: Source Identification, Information Extraction, Schema Alignment, Knowledge Graph Construction, and User Interface Development.

FoodKG, initially in the trig format, into nt triples format and, subsequently, to the KGTK format. Using Kypher (Chalupsky et al. 2021) we successfully extract granular nutrient and FOODON data from the broader dataset. Afterward, we use the Semantic Scholar Open Research Corpus (S2ORC) PDF parser (Lo et al. 2020) to extract information from academic articles.

Schema Alignment. We align ingredient data from three sources: nutritional content, the FOODON system, and water footprint information. First, we analyze ingredient names to extract main descriptors. This step simplifies complex names like 'vanilla-flavored soy yogurt' to 'soy yogurt'. Next, we turn graph relations into English sentences to compute sentence embeddings using Sentence Transformers. These embeddings serve as features for the model developed in the "Knowledge Graph Construction" phase.

Knowledge Graph Construction. We unify data from various sources using a common ingredient identifier. Addressing the challenge of missing entities in our knowledge base, we employ a neural network architecture: four dense layers followed by an output layer with a continuous ReLU activation function. Given an ingredient's embedding and its relation type, we predict the nutritional and water footprint value. It's optimized using the Adam optimizer and MSE loss. Our graph comprises of 20,778 nodes and 13 types of relationships, responding within an average latency of 1s.

User Interface Development. We design a user-centric interface via the Django framework. Users input their recipe ingredients, and via the Knowledge Graph we recommend alternatives with a lower water footprint, thus being more sustainable. We host this application on the AWS² platform.

²<https://aws.amazon.com/>

Demonstration

The demonstration video showcases the efficacy of our recommendation engine in promoting sustainable eating by emphasizing the water footprint. Upon entering recipe ingredients like butter cream, whipped cream, and chocolate fudge cookie, the system first discerns the 'parent ingredient' through the food ontology hierarchy. It then chooses 'candidate ingredients' from a similar tier, excluding those with a larger water footprint. Subsequently, ingredients are ranked, prioritizing those with the least water consumption. For instance, when users aim to reduce fat, the engine suggests alternatives. After opting for a recommended ingredient, the system visualizes the water footprint difference alongside any coincidental nutritional changes. In this example, the total water footprint significantly diminishes from 20,135m³/ton to 5,215m³/ton. While the primary focus is on water conservation, it's noteworthy that the fat content also decreases from 30.32g to 11.84g, underscoring the holistic benefits of our system.

Conclusion

In this work, we introduced a recommendation tool centered on the concept of water footprint, utilizing a knowledge graph pipeline to advocate for sustainable eating habits. Our tool prioritizes ingredient substitutions that not only enhance nutritional value but also significantly reduce the water footprint of recipes. The aim is to directly influence individual food choices in favor of sustainability and water conservation. Looking ahead, our future endeavors will involve conducting user studies, expanding our ingredient database with a focus on water footprint data, and refining the recommendation algorithm.

Acknowledgments

We thank Yuhua Wu and Shreya Padmanabhan for their invaluable contributions. This work was funded by the Defense Advanced Research Projects Agency with award HR00112220046.

References

- Chalupsky, H.; Szekely, P.; Ilievski, F.; Garijo, D.; and Shenoy, K. 2021. Creating and querying personalized versions of wikidata on a laptop. *Wikidata workshop*.
- da Silva, R. M. G. 2021. *Knowledge Graph-Based Recipe Recommendation System*. Ph.D. thesis, Universidade do Porto (Portugal).
- Dooley, D. M.; Griffiths, E. J.; Gosal, G. S.; Buttigieg, P. L.; Hoehndorf, R.; Lange, M. C.; Schriml, L. M.; Brinkman, F. S.; and Hsiao, W. W. 2018. FoodOn: a harmonized food ontology to increase global food traceability, quality control and data integration. *npj Science of Food*, 2(1): 23.
- Foley, J. A.; DeFries, R.; Asner, G. P.; Barford, C.; Bonan, G.; Carpenter, S. R.; Chapin, F. S.; Coe, M. T.; Daily, G. C.; Gibbs, H. K.; Helkowski, J. H.; Holloway, T.; Howard, E. A.; Kucharik, C. J.; Monfreda, C.; Patz, J. A.; Prentice, I. C.; Ramankutty, N.; and Snyder, P. K. 2005. Global Consequences of Land Use. *Science*, 309(5734): 570–574.
- Ge, M.; Elahi, M.; Fernaández-Tobías, I.; Ricci, F.; and Massimo, D. 2015. Using Tags and Latent Factors in a Food Recommender System. In *Proceedings of the 5th International Conference on Digital Health 2015*, DH '15, 105–112. New York, NY, USA: Association for Computing Machinery. ISBN 9781450334921.
- Harvey, M.; Ludwig, B.; and Elsweiler, D. 2013. You Are What You Eat: Learning User Tastes for Rating Prediction. In Kurland, O.; Lewenstein, M.; and Porat, E., eds., *String Processing and Information Retrieval*, 153–164. Cham: Springer International Publishing. ISBN 978-3-319-02432-5.
- Hausmann, S.; Seneviratne, O.; Chen, Y.; Ne’eman, Y.; Codella, J.; Chen, C.-H.; McGuinness, D. L.; and Zaki, M. J. 2019. FoodKG: A Semantics-Driven Knowledge Graph for Food Recommendation. In *The Semantic Web – ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part II*, 146–162. Berlin, Heidelberg: Springer-Verlag. ISBN 978-3-030-30795-0.
- Hoekstra, A. Y.; Chapagain, A. K.; Aldaya, M. M.; and Mekonnen, M. M. 2009. *Water footprint manual*. Water footprint network Enschede, The Netherlands.
- Ilievski, F.; Garijo, D.; Chalupsky, H.; Divvala, N. T.; Yao, Y.; Rogers, C.; Li, R.; Liu, J.; Singh, A.; Schwabe, D.; and Szekely, P. 2021. KGTK: A Toolkit for Large Knowledge Graph Manipulation and Analysis. arXiv:2006.00088.
- Liu, C.; Li, L.; Yao, X.; and Tang, L. 2019. A Survey of Recommendation Algorithms Based on Knowledge Graph Embedding. In *2019 IEEE International Conference on Computer Science and Educational Informatization (CSEI)*, 168–171.
- Lo, K.; Wang, L. L.; Neumann, M.; Kinney, R.; and Weld, D. 2020. S2ORC: The Semantic Scholar Open Research Corpus. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4969–4983. Online: Association for Computational Linguistics.
- Marrn, J.; Biswas, A.; Ofli, F.; Hynes, N.; Salvador, A.; Aytar, Y.; Weber, I.; and Torralba, A. 2021. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1): 187–203.
- Mekonnen, M. M.; and Fulton, J. 2018. The effect of diet changes and food loss reduction in reducing the water footprint of an average American. *Water international*, 43(6): 860–870.
- Mekonnen, M. M.; and Gerbens-Leenes, W. 2020. The Water Footprint of Global Food Production. *Water*, 12(10).
- Rijsberman, F. R. 2006. Water scarcity: fact or fiction? *Agricultural water management*, 80(1-3): 5–22.
- Springmann, M.; Clark, M.; Mason-D’Croz, D.; Wiebe, K.; Bodirsky, B. L.; Lassaletta, L.; De Vries, W.; Vermeulen, S. J.; Herrero, M.; Carlson, K. M.; et al. 2018. Options for keeping the food system within environmental limits. *Nature*, 562(7728): 519–525.
- van Pinxteren, Y.; Geleijnse, G.; and Kamsteeg, P. 2011. Deriving a Recipe Similarity Measure for Recommending Healthful Meals. In *Proceedings of the 16th International Conference on Intelligent User Interfaces, IUI '11*, 105–114. New York, NY, USA: Association for Computing Machinery. ISBN 9781450304191.
- Vermeulen, S. J.; Campbell, B. M.; and Ingram, J. S. 2012. Climate Change and Food Systems. *Annual Review of Environment and Resources*, 37(1): 195–222.
- Wada, Y.; van Beek, L. P. H.; van Kempen, C. M.; Reckman, J. W. T. M.; Vasak, S.; and Bierkens, M. F. P. 2010. Global depletion of groundwater resources. *Geophysical Research Letters*, 37(20).