

Leveraging Ontologies to Enhance Personalised and Sustainable Food Recommendations with LLMs

(Discussion Paper)

Ada Bagozi^{1,†}, Devis Bianchini¹, Michele Melchiori¹ and Anisa Rula¹

¹University of Brescia, Dept. of Information Engineering
Via Branze 38, 25123 - Brescia (Italy)

Abstract

Modern food recommendation systems aim to guide consumers toward sustainable and nutritionally balanced choices, encouraging healthier eating habits and addressing the increasing focus on food sustainability and waste reduction. Large Language Models (LLMs) have emerged as powerful tools for generating food recommendations due to their advanced natural language processing capabilities. However, delivering personalised and contextually accurate suggestions remains challenging due to the absence of a well-defined framework for representing healthy and sustainable food choices aligned with individual dietary and lifestyle needs. Ontologies provide a structured and semantically rich framework to address this gap. In this paper, we introduce a modular ontology designed to enhance LLMs-driven contextual understanding, enabling more accurate and personalised food recommendations. The ontology is built around competency questions drawn from a research project on sustainable and healthy food recommendations. We evaluated this approach by conducting experiments where ChatGPT-4 answered competency questions with and without ontology integration, followed by a user study to assess the responses. Initial results suggest that incorporating the ontology significantly improves the quality and relevance of the recommendations, with a positive impact on both search results and the quality of generated answers.

Keywords

Sustainable Food Recommendation, Consumer Empowerment, Multi-perspective Ontology Engineering, Large Language Models

1. Introduction

In the current food landscape, food recommendation systems help consumers make sustainable and nutritionally complete food choices, fostering healthy eating habits and addressing the growing interest in food sustainability and waste reduction [1]. While some approaches represent stakeholders' perspectives on sustainability and nutrition, consumers are a critical component of any food recommendation system [2]. Despite rising consumers' awareness, studies estimate that a substantial portion of food produced for human consumption is lost or wasted, with a significant share occurring at the household level [3]. Nutritional characteristics are another important dimension of food products, including composition, quality, and their impact on health [4]. Therefore, consumers need support in making food choices that are both sustainable and nutritionally complete, while adopting healthy eating habits. LLMs can be used for food recommendations by leveraging their ability to understand and process natural language [5]. However, providing personalised and contextually relevant suggestions remains challenging due to the integration and comprehension of domain-specific, heterogeneous sources, and the lack of a robust conceptualisation of healthy and sustainable food aligned with users' dietary and lifestyle preferences. Ontologies can address this issue by offering a structured and semantically rich framework for organising information. As shown in [6, 7], the combination of domain knowledge and LLMs shows promising preliminary results.

SEBD 2025: 33rd Symposium on Advanced Database Systems, June 16-19, 2025, Ischia, Italy

*Corresponding author.

✉ ada.bagozi (A. Bagozi); devis.bianchini@unibs.it (D. Bianchini); michele.melchiori@unibs.it (M. Melchiori); anisa.rula@unibs.it (A. Rula)

🆔 0000-0002-0193-6500 (A. Bagozi); 0000-0002-7709-3706 (D. Bianchini); 0000-0001-8649-4192 (M. Melchiori); 0000-0002-8046-7502 (A. Rula)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

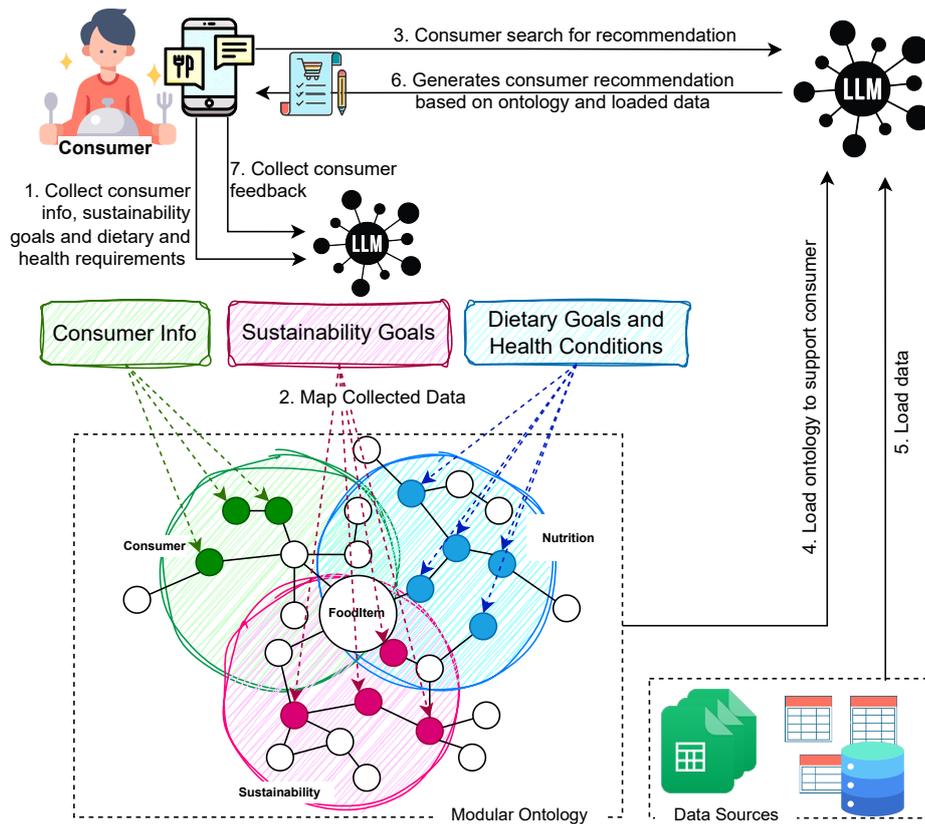


Figure 1: Overview on the combined use of LLMs and modular ontology for food recommendation.

In this paper, we propose a modular ontology to enhance the contextual knowledge of LLMs, enabling them to deliver personalised, contextually relevant food recommendations. We describe the process of building the ontological modules upon a set of competency questions, elicited within a research project focused on sustainable and healthy food recommendations. We discuss the ontology engineering process, its extensibility, and its current and future applications in modern food recommendation systems to address the following research questions: (i) How does the integration of a structured ontology with LLMs (specifically ChatGPT-4) improve the accuracy and relevance of food recommendations compared to using LLMs alone? (ii) What key dimensions and attributes must be included in a modular food ontology to support comprehensive and personalised food recommendations through LLMs? (iii) How is the use of an ontology with LLMs perceived by users in food recommendation systems, particularly concerning dietary restrictions and sustainability preferences? Our contribution combines ontology engineering with the advanced natural language processing capabilities of LLMs to address the challenges of integrating diverse data sources and providing tailored recommendations. Unlike existing approaches, our method ensures a comprehensive representation of consumers' preferences, dietary needs, and sustainability concerns, leading to more accurate and relevant food recommendations. A preliminary presentation of the approach has been provided in [8]. Here, we describe a series of experiments where ChatGPT-4 answered questions with and without ontology integration. The accuracy of the answers was evaluated through a user study. Preliminary experimental results demonstrate significant improvements in the quality and relevance of recommendations when the ontology is integrated.

The remainder of this paper is structured as follows: Section 2 provides the approach overview; Section 3 presents the methodology for requirements analysis and ontology construction; Section 4 describes the proposed validation process; Section 5 discusses the most recent work on the use of food ontologies to recommend food or to describe food products features to the final consumer, assuming a consumer-centric viewpoint; finally, in Section 6 we offer our conclusions and delineate possible future research directions.

2. Approach Overview

An overview of our approach is illustrated in Figure 1. The structure of the ontology is organised into three modules (consumer, nutrition, and sustainability modules) and facilitates the independent development and refinement of each module. The use of namespaces ensures that the ontology can be integrated seamlessly with other existing standardised vocabularies, and its application scope can be expanded with the addition of new concepts and properties as the domain evolves. For instance, new sustainability metrics or health parameters can be incorporated without significant restructuring. The integration of LLMs with the modular ontology provides an advanced method for interpreting user's inputs and generating relevant outputs. The proposed approach follows these steps:

1. **Collection of Consumers' Info:** data on consumers' dietary preferences, health conditions, and sustainability goals are gathered; this data forms the foundation for personalised recommendations.
2. **Mapping of Collected Data:** the LLM processes data previously collected and maps it into the modular ontology, ensuring structured and semantically rich representation.
3. **Food Recommendation:** consumers interact with the LLM to search for food recommendations, leveraging the ontology for contextually relevant results.
4. **Ontology Loading for Answers Generation:** the ontology is also exploited to support the LLM in answering the consumer's request, enhancing the accuracy and relevance of the generated responses.
5. **Data Loading:** the LLM loads data from various sources based on the ontology and user's inputs, ensuring that the recommendations are aligned with the user's preferences and goals.
6. **Generation of Recommendations:** the LLM generates personalised food recommendations based on the ontology and query results, detailing the nutritional benefits and sustainability impacts of each suggested food item.
7. **Feedback Collection:** the system collects consumers' feedback and interactions to improve future recommendations, ensuring continuous learning and adaptation.

The recommender system leverages data from various sources to retrieve food items and dietary plans that align with the user's preferences and goals. Based on the retrieved data, the LLM generates personalised recommendations, detailing the nutritional benefits and sustainability impacts of each suggested food item. This approach is essential for applications in several fields, including digital labelling [9], ontology-based data access [10], and LLM-based recommendation systems [11]. By integrating ontological context, the system enhances the relevance and quality of food recommendations, addressing the user's dietary needs and sustainability preferences effectively. This modular, ontology-driven framework ensures that the recommendations are comprehensive, accurate, and aligned with the latest research and consumers' feedback.

3. Requirements Analysis and Ontology Development

3.1. Requirements Analysis

To develop the food ontology, we conducted interviews with domain experts from various fields, including industrial engineering, dietetics, nutrition, agricultural economics, and food science. This process helped identify key knowledge aspects and use cases for the ontology construction. We collected information about experts' needs, to identify the main knowledge aspects they want to answer. Two primary use cases emerged. **UC1: Health and Nutritional Monitoring.** *Description:* this use case targets a consumer who wants to monitor and retrieve health parameters, dietary preferences, and physical activity levels to receive personalised nutritional and health advice. *Actors:* consumers and dietitians are involved in monitoring and managing health and well-being. *Flow:* consumers aim to receive food recommendations according to their health conditions and lifestyle preferences. Dietitians

will take care to provide data about micro-nutrition and macro-nutrition and correlations with physical conditions. **UC2: Sustainable Shopping and Consumption.** *Description:* this use case involves a consumer who wants to retrieve detailed information on the sustainability of food products, to make informed purchase decisions and to minimise waste. *Actors:* consumers, as well as different types of actors, are involved in the provisioning and exploitation of data on the sustainability of food products. *Flow:* consumers access the platform and browse detailed sustainability information on food products, allowing them to make informed purchasing decisions and minimise waste. The platform aggregates data from various stakeholders. This comprehensive data enables consumers to select products that align with their sustainability preferences. From these use cases, we derived a set of competency questions (CQs) to define the functional requirements of the ontology. These CQs cover aspects such as food impacts on health, nutritional content, dietary restrictions, sustainability metrics, and waste minimisation.

3.2. The Modular Food Ontology

The modular Food Ontology represents a significant advancement in the field of personalised food recommendation systems. The main structure of the ontology is shown in Figure 2. This comprehensive ontology is designed to address the complex interplay between consumers' preferences, nutritional requirements, and sustainability factors in food choice. By leveraging RDF and OWL technologies, the ontology provides a flexible and extensible structure capable of representing the multifaceted relationships inherent in the food domain. The reuse of existing classes and properties is done using the `rdfs:subClassOf` or `owl:equivalentClass` construct for classes and the `rdfs:subPropertyOf` construct for properties.

Consumer Module. This module is designed to capture a comprehensive view of a consumer's lifestyle, encompassing demographic information, activities, dietary choices, health conditions, and shopping habits. By integrating various aspects of daily life, it provides a holistic understanding of consumer's behaviour. As shown in Figure 2, this module comprises 7 classes identified by the yellow colour. Central to this module is the `Consumer` class, which extends `schema.org's Person` class to capture detailed personal and lifestyle-related information. It links shopping activities to food choices and health conditions through classes such as `HealthConditions` and `RestrictedDiet`. Sustainability is another relevant dimension, captured by the `Sustainability` class. This reflects consumers' preferences for sustainable practices, highlighting the growing importance of environmental considerations in consumers' behaviour.

Nutrition Module. The nutrition module represents the intricate relationships between consumers, their dietary plans, and the various health and personal factors influencing their nutritional requirements. As shown in Figure 2, this module comprises 8 classes identified by the green and the lilac colour. At the core of this module is the `HealthConditions` class, which captures various health-related aspects of the consumer. It includes specific conditions such as food allergies, intolerances and food-related pathologies. These conditions play a crucial role in shaping consumer's dietary requirements. To this purpose, the `RestrictedDiet` class models various dietary restrictions and preferences, while the `FoodItem` class represents individual food items relevant to `RestrictedDiet` and `HealthConditions`.

Sustainability Module. The sustainability module captures a comprehensive view of sustainability within the food supply chain, emphasizing the consumer's role in shaping sustainable practices. It integrates environmental, social, and economic dimensions to reflect their impacts on various stakeholders such as `Farmer`, `Producer`, and `Retailer`. As shown in Figure 2, this module comprises 7 classes identified by the blue colour. At the core of this module is the `Sustainability` class, that is specialised into 4 subclasses: `EnvironmentalSustainability`, `SocialSustainability` and `EconomicSustainability` to model the impact of food production and consumption on various stakeholders. By leveraging existing ontologies such as `Agronomy Ontology (AgrO)` and `Environmental Ontology (ENVO)`, this module facilitates the analysis of sustainable practices throughout the food supply chain.

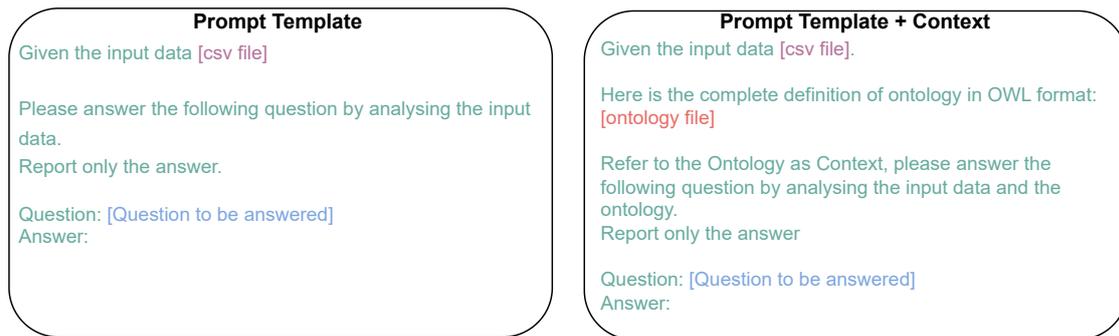


Figure 3: Prompt template and prompt template with the ontology.

participants couldn't distinguish between the two types. They rated the solutions according to five criteria: *Accuracy* (correctness of the answer), *Completeness* (coverage of the answer), *Clarity and Structure* (how much understandable is the answer), *Depth and Detail* (level of explanation of the answer), and *Practicality and Usefulness* (actionable value), with scores ranging from 1 (low) to 5 (high).

Datasets. The datasets utilised to generate the responses are retrieved online or automatically generated¹, as reported in the following: (i) **food.csv** contains comprehensive data about different types of food (CORGIS); (ii) **nutrition.csv** provides nutritional values for common foods and products (Kaggle); (iii) **generic-food.csv** includes general information about various food items (GitHub); (iv) **food-production.csv** details the environmental impact of food production (Kaggle); (v) **customer_activities.csv** contains data about customer's activities (generated); (vi) **detailed-dietary-requirements.csv** provides detailed dietary requirements for various consumers (generated); (vii) **detailed-sustainability-preferences.csv** lists detailed sustainability preferences of consumers (generated); (viii) **purchase_data.csv** contains data about past purchases made by consumers (generated); (ix) **complete.owl** contains the OWL representations of the complete ontology loaded and maintained in the **GitHub repository** (the proposed ontology).

Prompt templates design. We propose structured prompt templates for food recommendations by creating a template with placeholders for input data and ontology (see Figure 3). There are two types of templates: one with ontology and one without. The ontology-based template uses intensional definitions to provide contextual knowledge, improving the understanding of the model and the relevance of its recommendations. These templates also include placeholders for competency questions derived from users' stories and use cases, covering both simple and complex queries. We used GPT-4, hosted on ChatGPT-Plus, with in-context learning to better align the model with specific tasks. All experimental data and results have been loaded in the **GitHub repository**.

4.2. Preliminary Evaluation

Figure 4(a) illustrates the overall evaluation given by the experiment participants across the two solutions: "Data and Ontology" and "Only Data." The data collected across nine questions (Q1 to Q9) indicates a clear preference for the integrated solution, with more than 70% of users favouring "Data and Ontology" for most questions. This preference underscores the significant enhancement in recommendation accuracy when contextual information from an ontology is utilised. The bar chart in Figure 4(a) clearly depicts the performance gap, showing that the system incorporating context consistently outperforms the data-only approach across all questions. For instance, Q1 and Q7 show notable improvements with context, scoring 71% and 79%, respectively, compared to 29% and 21% without it. Particularly striking are the results for Q4, Q5, and Q9, where the context-enhanced system

¹During dataset preparation, the authors employed ChatGPT-4 to generate synthetic data. All outputs were subsequently reviewed and edited by the authors, who assume full responsibility for the final content of the publication.

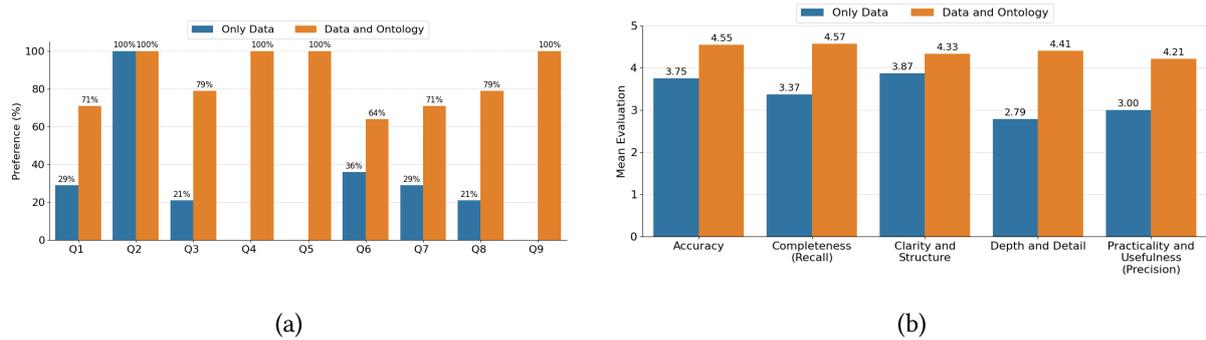


Figure 4: Percentage of preferences expressed for each solution Figure 4a and overall mean evaluation of parameters Figure 4b

was preferred by all participants (100%), while the data-only approach received no preference (0%). Moreover, Q5 highlighted a notable difference between the two solutions. The context-enhanced response provided comprehensive recommendations for Jason Brown’s weekly shopping, aligning with his dietary preferences and sustainability goals. In contrast, the data-only solution failed to identify any suitable items due to limitations in the datasets. The language model clarified that the context provided by the ontology includes a more comprehensive understanding of food categories, sustainability attributes, and dietary requirements, allowing for broader interpretations beyond the raw data in the CSV files. This context helps to bridge the gaps where direct data filtering falls short, enabling informed recommendations even in the absence of matching items in the datasets. Question Q6 also exhibits significant gains, with context-based performance at 64% versus 36% for data alone. For Q2, both solutions were equally preferred (100% for both), indicating that for questions focusing on technical information like macro-nutrients and micro-nutrients in salmon, the integration of ontology does not necessarily change users’ preference as long as the provided data is accurate and complete. Overall, these findings highlight the critical importance of integrating contextual knowledge to enhance the system capability in handling complex queries, demonstrating that context not only improves accuracy, but is sometimes essential for success.

Figure 4(b) presents the overall mean evaluation scores comparing the two solutions. The bar chart reveals a substantial improvement in users’ satisfaction with context integration, reflected across the five evaluation criteria: *Accuracy*, *Clarity and Structure*, *Completeness (Recall)*, *Depth and Detail*, and *Practicality and Usefulness*. The context-integrated system consistently scores higher, with notable differences such as a mean accuracy of 4.55 compared to 3.75 for the data-only system, and a remarkable 4.57 for completeness versus 3.37. Users found the context-enhanced responses of the system clearer and more structured (4.33 vs. 3.87), significantly richer in depth and detail (4.41 vs. 2.79), and more practical and useful (4.21 vs. 3.00). These results underscore the critical role of contextual knowledge in enhancing the overall performance of the system, making recommendations more accurate, comprehensive, and user-friendly.

5. Related Work

Food recommendation systems have seen significant advancements through the integration of ontologies, knowledge graphs, and LLMs. These systems aim to promote healthy eating habits, sustainable food choices and reduce waste. Our work builds upon and extends several key areas of research in this field. **Ontology-Based Approaches:** Comprehensive food ontologies such as FoodOn [12] and PO2/TransformON [13] have provided structured frameworks for food domain modelling and sustainability in the agri-food supply chain. These ontologies offer rich representations of food systems but often lack direct integration with advanced AI techniques for personalised recommendations. Ontology-based systems have shown potential in providing tailored nutrition advice and adapting to specific users’

groups, demonstrating their versatility in food recommendation contexts. **Knowledge Graphs in Food Recommendation:** Knowledge graphs have emerged as powerful tools for representing complex food-related information. Haussmann et al. [14] introduced FoodKG, a semantically rich food knowledge graph that integrates data from various sources to support food recommendation tasks. Recent work by Simsek-Senel et al. [2] further demonstrates the potential of knowledge graphs in modelling sustainable food production. **Integration with AI techniques:** The integration of ontologies and knowledge graphs with LLMs represents a promising frontier in recommendation systems. Pan et al. [6] provided a roadmap for unifying LLMs and knowledge graphs, while Allemang and Sequeda [15] demonstrated how ontologies can increase LLM accuracy for question answering. Sequeda et al. [7] showed that LLM-powered systems leveraging knowledge graph representations can significantly improve accuracy in specialized domains, returning up to three times more accurate results than LLMs without such knowledge structures. **Consumer-centric and sustainable approaches:** Consumer-focused studies, such as Botos et al. [9], highlight the growing demand for easily accessible, personalised food-related data. Sustainability and waste reduction have emerged as crucial aspects, with Hermanussen and Loy's [3] meta-analysis underscoring the importance of addressing food waste at the consumer's level. Despite these advancements, there remains a gap in comprehensively integrating nutritional, sustainability, and personal preference factors within a single system. Moreover, the potential of combining advanced LLMs with domain-specific ontologies for food recommendations has been largely unexplored. Our research addresses these gaps by proposing a modular ontology that enhances the contextual knowledge of LLMs, aiming to improve the accuracy and relevance of food recommendations while promoting sustainability and healthy eating habits.

6. Concluding Remarks

In this paper, we presented a modular ontology to enhance LLMs in delivering personalised, contextually relevant food recommendations. By integrating structured ontology with LLMs, we address the limitations of using LLMs alone, improving both the accuracy and relevance of recommendations. Our experiments with ChatGPT-4 demonstrated that ontology integration significantly boosts the quality and contextual relevance of responses.

Future work will focus on integrating Retrieval-Augmented Generation (RAG) techniques to combine structured and unstructured data, further enhancing recommendation quality. Additionally, we plan to conduct larger-scale experiments with a more diverse user base, including the collection of real customers' data to better capture preferences and behaviours. Continuous refinement of the ontology, with feedback from these experiments, will expand its scope to cover additional subcategories of food sustainability and nutrition. These improvements will contribute to building more sustainable and health-conscious food systems globally.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

- [1] A. Konys, An ontology-based knowledge modelling for a sustainability assessment domain, *Sustainability* 10 (2018) 300.
- [2] G. Simsek-Senel, H. Rijgersberg, B. Öztürk, J. Weits, A. Fensel, I-know-foo: Interlinking and creating knowledge graphs for near-zero co2 emission diets and sustainable food production, in: *AI, Data, and Digitalization, 2024*, pp. 106–119.
- [3] H. Hermanussen, J.-P. Loy, Household food waste: A meta-analysis, *Environmental Challenges* 14 (2024) 100809.

- [4] A. Brena-Melendez, L. E. Garcia-Amezquita, A. Liceaga, C. Pascacio-Villafán, V. Tejada-Ortigoza, Novel food ingredients: Evaluation of commercial processing conditions on nutritional and technological properties of edible cricket (*acheta domesticus*) and its derived parts, *Innov. Food Sci. Emerg. Technol.* (2024).
- [5] Y. Deldjoo, F. Nazary, A. Ramisa, J. J. McAuley, G. Pellegrini, A. Bellogín, T. D. Noia, A review of modern fashion recommender systems, *ACM Comput. Surv.* (2024).
- [6] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, X. Wu, Unifying large language models and knowledge graphs: A roadmap, *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [7] J. Sequeda, D. Allemang, B. Jacob, A benchmark to understand the role of knowledge graphs on large language model's accuracy for question answering on enterprise SQL databases, in: *Workshop GRADES and (NDA)*, ACM, 2024.
- [8] A. Bagozi, D. Bianchini, M. Melchiori, A. Rula, Enhancing llms contextual knowledge with ontologies for personalised food recommendation, in: *The Web Information Systems Engineering (WISE 2025)*, Doha, Qatar, 2024, pp. 273–283.
- [9] S. Botos, M. Tóth, R. Szilágyi, Improving food consciousness - opportunities of smartphone apps to access food information, *J. Agric. Inform* 12 (2022).
- [10] S. Hoseini, J. Theissen-Lipp, C. Quix, A survey on semantic data management as intersection of ontology-based data access, semantic modeling and data lakes, *JWS* 81 (2024).
- [11] Y. Hou, J. Zhang, Z. Lin, H. Lu, R. Xie, J. McAuley, W. X. Zhao, Large language models are zero-shot rankers for recommender systems, in: *Advances in Information Retrieval*, 2024, pp. 364–381.
- [12] D. M. Dooley, E. J. Griffiths, G. S. Gosal, P. L. Buttigieg, R. Hoehndorf, M. C. Lange, L. M. Schriml, F. S. Brinkman, W. W. Hsiao, Foodon: a harmonized food ontology to increase global food traceability, quality control and data integration, *npj Science of Food* 2 (2018) 23.
- [13] M. Weber, P. Buche, L. Ibanescu, S. Dervaux, Po2/transformon: A new domain ontology for integrating food, feed, bio-products and waste in a circular and sustainable approach, *CEUR Workshop Proceedings* 3637 (2023) 1–12.
- [14] S. Hausmann, O. Seneviratne, Y. Chen, Y. Ne'eman, J. Codella, C.-H. Chen, D. L. McGuinness, M. J. Zaki, 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* 18, Springer, 2019, pp. 146–162.
- [15] D. Allemang, J. Sequeda, Increasing the LLM accuracy for question answering: Ontologies to the rescue!, *CoRR* (2024).