NutriWell: an Explainable Ontology-Based FoodAI Service for Nutrition and Health Management

Berardina De Carolis[1 [0000−0002−2689−137X]], Davide Lofrese[1], Davide Di Pierro[1 [0000−0002−8081−3292]], and Stefano Ferilli[1 [0000−0003−1118−0601]]

Università degli Studi di Bari Aldo Moro, Bari, Italy
name.surname@uniba.it

Abstract. Non-communicable diseases (NCDs) like hypertension, diabetes, osteoporosis, and cancer constitute 80% of the disease burden in European countries, affecting a significant portion of the working-age population. Addressing these numbers requires a strong effort in prevention and management. Nutrition is crucial not only for chronic conditions but also for non-chronic medical needs such as pregnancy, allergies, and intolerances. Artificial Intelligence (AI), especially when integrated with chatbots or social robots, now plays a pivotal role in assisting users with NCD prevention and management, as well as dietary needs. This manuscript introduces *NutriWell*, a framework leveraging AI and the *GraphBRAIN* technology for intelligent knowledge retrieval in nutrition and health management. *NutriWell* informs users about meal suitability based on their nutritional requirements, utilizing explanations that combine feature data and user preferences. Italian websites such as *GialloZafferano* and *AlimentiNUTrizione* provide extensive catalogs of European meals, including ingredients, allergens, and dietary specifics. The contribution of this work is the construction of a personalized diet assistant by utilizing datasets extracted from these websites that, as far as we know, have never been used for these tasks. A key contribution is an API that retrieves graph-based information integrated with an ontology specifying relational constraints. The ontology design, derived from existing frameworks and enhanced to integrate food impacts on disorders, allows for the calculation of meal impact scores tailored to user needs and preferences.

Keywords: Graph Database · Ontology · Health.

1 Introduction

Non-communicable diseases (NCDs), also known as chronic diseases, are diseases that are not transmissible from one person to another and include hypertension, diabetes, osteoporosis, and cancer. NCDs are responsible for 80% of the disease burden in European countries [23] and about one-quarter of the working-age population (23.5%) is affected by at least one chronic disease [22]. Furthermore, multiple chronic conditions (MCCs) often occur at the same time, which is very common in people aged more than 65 years old (65%) and up to 85% in elderly people (aged 85+) [17]. In the next years, these figures are predicted to increase. For example, diabetes diagnoses in Europe have almost tripled over the last two decades, from about 22 million in 2000 to 61 million in 2021, and 2045 expects a further increase of 13% [7]. Genetic and environmental factors contribute to the risk of developing one or more NCDs, but modifiable behavioural risk factors play an essential role. In particular, tobacco use, lack of physical activity, and unhealthy diet increase the risk of NCDs. Indeed, the rise in chronic medical conditions can be linked to diets becoming more and more characterized by high levels of fats and sugars, but lacking in fresh fruits and vegetables [4]. Moreover, about two million deaths each year are attributed to excess salt and sodium intake [8].

Hence, a healthy diet and lifestyle can help in preventing the onset of such diseases. For instance, a diet that is low in fats, refined carbohydrates, and sugar-sweetened beverages, but rich in fruits and vegetables reduces the risk of cardiovascular diseases [24]. In addition, following a healthy diet is useful not only in the prevention of an NCD but also in the treatment and control of an already-developed disease. An example is the DASH (Dietary Approaches to Stop Hypertension) diet which helps reduce blood pressure in hypertensive patients.

Improving eating habits is important for other non-chronic medical conditions as well, such as pregnancy where a high supply of nutrients is necessary, but limiting sugars, salt and saturated fats. Furthermore, people affected with food allergies or intolerances have to be careful in what they eat to avoid the unintentional intake of unsafe food containing allergens that could trigger life-threatening reactions.

To make one's dietary habits better, food and nutrition literacy became very important [18]. Specifically, nutrition literacy is the ability to obtain, understand, and apply nutritional information to make healthy decisions about food consumption. A nutrition-literate person is aware of the food nutrients and how they affect health and, as a consequence, he/she can exercise good judgment in choosing what to eat.

Technologies can help in increasing nutrition literacy and gaining a greater level of nutritional awareness by providing hardware and software tools to support people in managing their nutrition and food intake. For example, digital food diaries can be used to log meals and monitor calories and macronutrients (i.e., carbohydrates, proteins and fats) to keep a balanced diet. Similarly, virtual coaches in the form of chatbots can provide users with personalized diet recommendations based on age, gender and health conditions.

Due to the nature of the problem and the necessity of making people understand what modern AI systems output, explanations for the results need to be provided. We emphasize the personalization factor in this approach since dietary habits, diseases and intolerances make each one of us different.

In this work, we propose a new framework for evaluating how appropriate is a meal for an end-user, taking into account user intolerances, diseases, preferences (e.g. vegetarianism), food ingredients, calories and more. The dataset is constructed by combining information available in some of the most comprehensive food websites, such as *GialloZafferano* and *AlimentiNUTrizione*. The combined dataset has been translated into a graph formalism (specifically Labelled Property Graph [1]) which can be accessed through an API. The graph is empowered with an ontology defining entity and relationship constraints. The ontology also acts as a scheme for the graph and has been designed starting from existing schemes and connecting information about foods with their impact on disorders that is, to the best of our knowledge, an underrepresented way of formalizing schemes in the food domain. The graph and the ontology follow the *GraphBRAIN* technology, next described. This manuscript is structured as follows: Section 2 presents existing solutions and comparisons with our framework, Section 3 describes data gathering and storage, Section 4 describes the *GraphBRAIN* framework and the ontology design process, Section 5 describes the overall *NutriWell* framework, Section 6 describes the explanatory process and the evaluation, Section 7 clarifies possible ethical concerns, and finally Section 8 concludes the work and mentions future scenarios.

2 Related Work

The food sector has largely benefited from Artificial Intelligence. The applications are countless [13], but they can be mainly divided into two categories: the control and tracking, from the harvesting to the consumption, and the effects of the aliments on our bodies, more often than not proposing recommendations [21]. One of the first attempts to advertise meals based on the user needs was by Snae et al. [19]. It was based on Korean cuisine and represented a filtering system for diseases. From this earlier stage, we fuzzify the impact by assigning a certain degree of damage (not just a binary choice) and provide explanations by combining ingredients' features.

A more sophisticated recommender system by Toledo et al. [20] has personalization as its main strength and also considers user preferences. One of the main differences with our approach is that we are capable of providing explanations for informing the user about the goodness of a choice. This is due to the use of interpretable graph models (graphs) and symbolic conceptualizations (ontologies).

Regarding the pure ontology designing phase, the attempt to unify food ontologies has been provided by Popovski et al. [14]. However, the unified conceptualization does not include peculiarities details for expressing the impact of the food and ingredients on allergies and/or diseases. In the work, the alignment took into consideration several existing ontologies such as FoodWiki [3], Open Food Facts [16] and others. Nonetheless, these resources cannot be compared with the richness of the Italian websites we are considering here. On the other side, our solution lacks the presence of foreign cuisines.

In the context of food production tracking for sustainability and smart cities, Kamel et al. [12] integrated heterogeneous sources to combine information at different granularity and steps for a holistic production line of food.

In dealing with graphs, Qin et al. [15] constructed a Chinese cuisine knowledge graph with a query answering system retrieving resources gathered from the web. In this work, the ontology is derived from data rather than manually constructed and refined for the task, also because the system was not supposed to be used for a specific use case. Again, the lack of correlation between food and illnesses, or the correlation with the ingredients' features is common.

Without using conceptualization on a graph, as often happened with the LPG model, Bajaj et al. [2] introduced the use of Neo4j for the storage and recommendation of food. This represented a preliminary idea of our work in which the performance can be appreciated and justifies the ongoing research in this field with that representation.

Taking everything into account, many works aimed to recommend personalized food based on diseases/allergies, some of them are not explainable or limited in their interpretation since they do not make use of interpretable data structures or ontologies. The conceptualization of food barely takes into account intolerances and the impact of each ingredient on them. Datasets vary in nationality and there have been attempts to merge them but, as far as we are concerned, nobody took into account two of the most popular Italian websites.

3 Dataset

The first step was to realize a new dataset, that is the result of the combination of existing datasets available in the form of websites. As far as we are concerned, it is the first time *GialloZafferano* and *AlimentiNUTrizione* have been employed for developing an AI-based diet assistant. The main limitation of this dataset is that it is in Italian but it may represent a relevant resource in the field for the community. One of the contributions consists of the collection of these web-based data (through scraping), translation, and dissemination. Diseases, allergies, and intolerances have been manually selected, considering their relation to food, and manually listed by experts who examined the dataset. Afterwards, foods, beverages and ingredients have been provided with an impact on the diseases, which can be positive, neutral, or negative. This assumption holds if the quantities respect the traditional recipe more or less. For instance, an exceeding amount of salt harms every person regardless of his/her physical state.

3.1 Data Gathering

The majority of data on foods comes from the two aforementioned websites. The first one is the most visited recipes website in Italy and contains more than 4,000 recipes. It has been chosen as a data source both for its numerosity and for the presence of nutritional information. Specifically, the nutrients available in *GialloZafferano* are eight: carbohydrates, proteins, fats, sugars, saturated fats, fibers, cholesterol, and sodium. The second website, on the other hand, is curated by the Italian Food and Nutrition Research Center and includes the Food Composition Tables, a collection of composition data for about 900 basic foods. We also included some other sources like *Cookaround* and *leCalorie* to gather a few regional recipes. In summary, *GialloZafferano*, along with other recipe websites, is the data source for nutritional facts of complex recipes (e.g., cheesecake), whereas *AlimentiNUTrizione* is the source for nutritional facts of basic foods (e.g., apple). Fortunately, ingredients in *GialloZafferano* did not need any entity recognition process since no synonyms were used throughout the websites. On the other hand, ingredients in *AlimentiNUTrizione* were manually linked to those available in *GialloZafferano*.

3.2 Data Labelling

After a data cleaning process to remove invalid values and/or redundancies, the labelling step took place. The complete set of the ingredients contained in the recipes has been taken into account to label each ingredient with a category (e.g., fruit, meat, sauce). For this purpose, a set of categories has been defined based on the classes in the HeLiS ontology [5]. In particular, this ontology provides a representation of the food and physical activity domains and has been used in a real-world system for promoting healthy lifestyles in workplaces. Among the food-related concepts included in the ontology, the ingredients are broken down into 137 categories organized in a hierarchical taxonomy. These concepts are represented with a high level of granularity, nonetheless, for this task, a subset of these classes has been considered discarding too specific classes. For instance, the **Red Meat** class has been included, but specific types of red meat have been neglected (e.g., sheep, bovine, goat, and pork red meat). The total classes selected are 35 and the resulting hierarchical structure is shown in Figure 1.

After defining the categories to use, the ingredients have been manually assigned to the most specific category. For example, **Sliced Beef** has been assigned the **Red Meat** category, instead of the more generic **Meat** category.

Labelling ingredients allows assigning diet and health labels to foods and beverages. More specifically, as in Edaman [6], it has been chosen to supplement foods and beverages with a set of diet and health labels to provide information on nutrient-level and ingredient-level aspects of the foods.

The set of diet and health labels used in this work is based on a subset of the labels provided by Edamam. In particular, 12 labels have been defined (5 diet labels and 7 health labels) and a short description of each is reported in Table 1.

In total, 21 conditions have been chosen, with 12 chronic diseases and health conditions (acne, cancer, diverticulitis, high cholesterol, hypertension, irritable bowel syndrome, migraine, non-alcoholic fatty liver disease, obesity, osteoporosis, pregnancy, type 2 diabetes), 6 food allergies (egg, fish, mollusc, peanut, shellfish, tree nut), and 3 food intolerances (fructose, gluten, lactose).

Following the selection of diseases, a range of values has been defined to express the positive or negative effect of a food, beverage, or ingredient on a

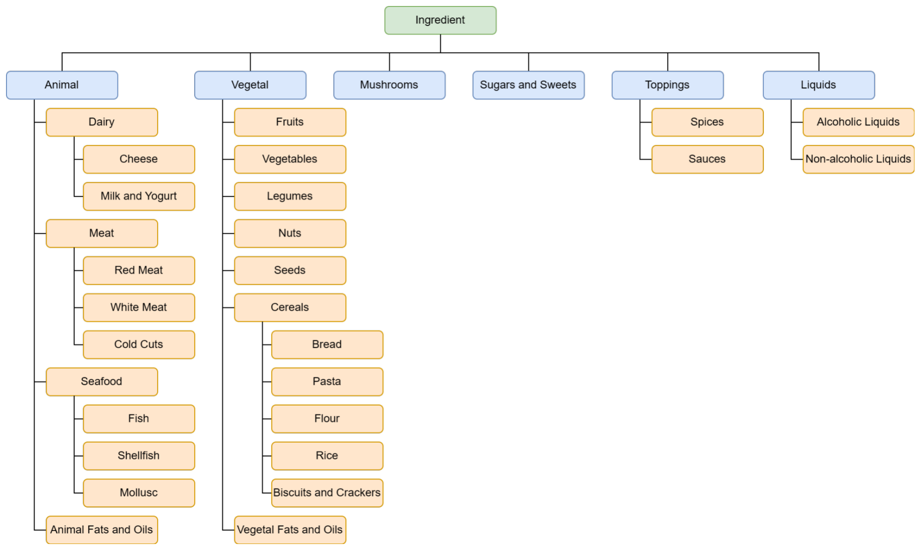

Fig. 1. Ingredients' ontology.

disease, allergy or intolerance. In particular, it has been chosen to assign values in the closed interval $[-1, 1]$ with the following interpretation:

– -1: the food, beverage or ingredient is bad and should be avoided as it worsens or increases the risk of the condition
– 0: the food, beverage or ingredient has a neutral or no impact on the condition (this value is also used when there is no information about the relationship on the disease)
– $+1$: the food, beverage or ingredient has a helpful effect as it improves or prevents the condition

Note that, for generic diseases and medical conditions, values in-between -1 and $+1$ are possible. For instance, a value of -0.5 represents a food, beverage or ingredient that is moderately bad and should be consumed less, but not avoided. The choice of this range has been suggested by experts in the domain, who claimed that a higher level of granularity requires a more fine-grained user knowledge, and hence is not general.

Conversely, for allergies and intolerances, only the three values -1, 0 and $+1$ are possible without in-between values, that is:

– -1: the food, beverage or ingredient cannot be eaten as it contains the allergen
– 0: the food, beverage or ingredient is a little risky and might contain traces of the allergen (this value is also used when there is no information about the relationship on the allergy/intolerance)
– $+1$: the food, beverage, or ingredient is safe to eat as it does not contain the allergen

Table 1. Diet and health labels

Diet	balanced	Balanced ratio of carbohydrates, proteins and fats. Specifically: • % carbohydrates from 45 to 65 • % proteins from 10 to 35 • % fats from 20 to 35
	high-protein	% proteins > 50
	low-carbs	% carbohydrates < 20
	low-fat	% fats < 15
	low-sugar	# grams of sugars for serving < 4
Health	alcohol-free	No alcoholic ingredients
	pork-free	No pork meat ingredients
	red-meat-free	No red meat ingredients
	no-oil-added	No oil added in the food
	pescatarian	No meat ingredients
	vegan	No animal ingredients
	vegetarian	No meat or fish ingredients

First, the ingredient-disease pairs have been manually labelled according to knowledge about the interaction between ingredients and diseases. For instance, the ingredient "Salt" has been assigned the value -1 for hypertension disease as it is one of the ingredients directly linked to high blood pressure and, as such, it should be avoided. For allergies and intolerances, the labelling has been carried out checking, for each ingredient, whether it contains the allergen. For example, the ingredient "Dark chocolate" has been assigned the rating 0 for the egg food allergy, as it might contain the allergen. Next, food-disease pairs have been automatically labelled considering the ingredients in the foods and the previously labelled ingredient-disease pairs. Specifically, the assigned value for a food-disease pair is the average of the ingredient-disease values for each ingredient present in the food. For example, if there is a food containing four ingredients with -1, -1, $+1$ and 0 as ingredient-disease ratings for a specific disease, the food-disease pair will be labelled with -0.25, that is the average of the four ingredient-disease values. This is a safe practical simplification in general, needed to make the process of labelling each recipe feasible. However, this does not apply to allergies where the assigned value for a food-allergy pair is the the minimum of the ingredient-disease values, rather than their average. Consequently, if the food has even one ingredient containing the allergen, the food-allergy pair will be labeled with -1 as it is not safe to eat.

The final dataset is made of 5,239 foods and beverages (4,840 recipes, and 399 basic foods) along with their nutritional facts. In addition, the dataset includes 1,606 ingredients divided into 35 categories. Each food is labelled with a set of diet and health labels chosen among 12 possible labels. Furthermore, the dataset features information on the interaction of foods and ingredients with 21 diseases, allergies and intolerances.

4 GraphBRAIN and Ontology Alignment

For the storage of the dataset, many alternatives are available currently. Due to the relational structure of the data (e.g. connections among foods and ingredients, and ingredients with diseases and allergies), and the need for both competitive performance and interpretability, we opted for a graph-based representation, specifically the Labelled Property Graph [1] supported by Neo4j. While providing efficient query-answering techniques and path interpretability, graph databases lack a general structure and schema, which is a requisite for traditional relational databases [11], making the interpretation of labels blurry when evaluating the different results. Following this need, in *GraphBRAIN* (GB) [9, 10] we created a framework to deal with LPG data with a manually-defined upper schema on them. Schemes are the abstraction of nodes and arcs in the graph, representing the following (main) concepts:

– **entity**: set of nodes sharing the same label.
– **relationship**: arc between nodes.
– **attribute**: property of a node (resp. arc) in the LPG.
– **hierarchy**: a label being a specification of another one.

As it can be noticed, schemes are interpreted as ontologies. In the GB setting, data are fully compliant with schemes and no exceptions are allowed. For this purpose, part of the contribution consists in the developing of a new scheme for the food domain collecting structural elements of foods and ingredients, and connecting them with the impact on preferences, allergies and diseases.

A preliminary food conceptualization was already available in GB, but it has been expanded to fulfil the requirements of this work. The main elements of the dataset are these three concepts:

– **Aliment**: a generic aliment, specialized in Food and Beverage.
– **Ingredient**: a generic ingredient, specialized in 35 classes (see Figure 1).
– **Disease**: a generic disease, specialized in Allergy and Intolerance.

The whole conceptualization (visualized as a graph) is shown in Figure 2.

GB is available in the form of API to allow end-users to build their graph-based application.

5 NutriWell

NutriWell, the proposed framework, provides intelligent access to nutritional information about foods and beverages (with details on macronutrients, micronutrients and ingredients) but also gives insights on the impact of food on diseases, allergies, and intolerances. This information is accessible using a service designed and developed as a web API following REST architecture principles.

Resources exposed by the API have been identified based on the concepts represented as classes and relationships in the ontology. Specifically, the three concepts defined as top-level classes in the ontology, like **Aliment**, **Ingredient**

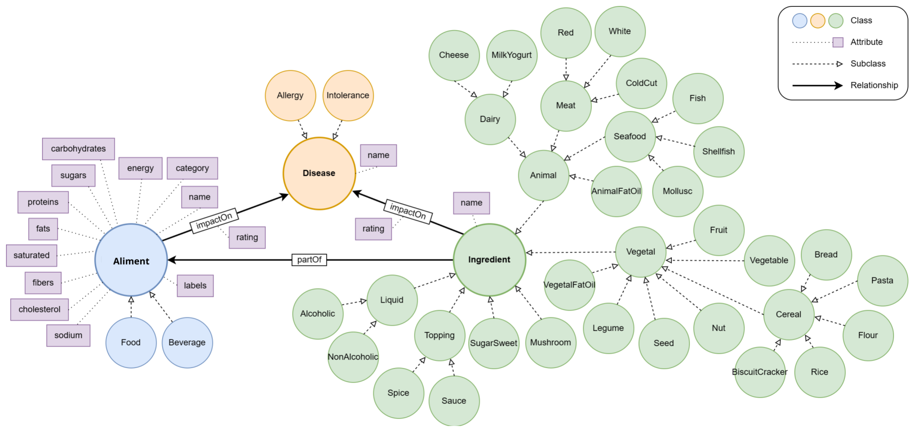

Fig. 2. GB Food Ontology

and **Disease**, give rise to two granularities of REST resources: collection resources and singleton resources. The former represents groups of homogeneous items, whereas the latter represents specific items within a collection. For example, the **Aliment** concept results in a collection resource, including all the foods and beverages, and multiple singleton resources, one for each specific food or beverage. The same applies to the **Ingredient** and **Disease** concepts, where a collection groups together the items, which can also be accessed individually.

Relationships among the concepts are represented as nested sub-collection resources. In particular, the *partOf* relationship specifying the ingredients contained in a food is represented as a sub-collection nested within the corresponding singleton food resource. In the opposite direction, the set of foods containing an ingredient is modelled as a sub-collection within the singleton ingredient resource. Similarly, the impact scores of a specific food (ingredient) on diseases are nested within the singleton food (ingredient) resources and, in the opposite direction, the impacts of foods (ingredients) on a specific disease are nested within the singleton disease resource.

Each singleton resource is described with a set of attributes that match the attributes of the respective classes in the ontology. Furthermore, two additional attributes are included in each singleton resource: *id* (the unique numeric identifier of the item) and *type* (the most specific ontology class assigned to the item).

Resources are represented in JSON format and they are accessed via URLs organized into a hierarchy. More specifically, collection resources have URLs based on plural nouns (*/foods*, */ingredients*, */diseases*), whereas singleton resources are assigned a unique numeric identifier and accessed through their parent collection resource (*/foods/{food-id}*, */ingredients/{ingredient-id}*, */diseases/{disease-id}*). Instead, nested sub-collection resources are accessed through their parent singleton resource (e.g., */foods/{food-id}/ingredients*, */foods/{food-id}/diseases*).

Operations on resources are defined in terms of standard HTTP methods identifying the type of operation to carry out. *NutriWell* uses five HTTP methods with their standard semantics, namely: GET, POST, PUT, PATCH and DELETE. Hence, the four traditional CRUD operations can be executed on resources. This allows data to change over time by adding new foods, ingredients and diseases.

The API also implements filtering features allowing clients to obtain subsets of large collection resources. In particular, one or more filters can be specified to retrieve a subset of items satisfying some constraints. For instance, the *name* filter parameter can be used to retrieve foods and beverages matching a string according to a similarity measure (namely, the Sørensen-Dice similarity measure).

NutriWell is implemented in Python leveraging *Litestar*, an open-source framework focused on building APIs. The system architecture is outlined in Figure 3. The server consists of several modules with well-defined responsibilities that cooperate to process requests. In short, a request coming from the client is received from the router and passed to the authorization middleware to check the API key validity. Then, the request reaches the controller that handles it using the service layer which, in turn, accesses the database through the repository layer. Note that the repository does not directly interact with the Neo4j database, but instead, it uses *GraphBRAIN* functionalities through its API.

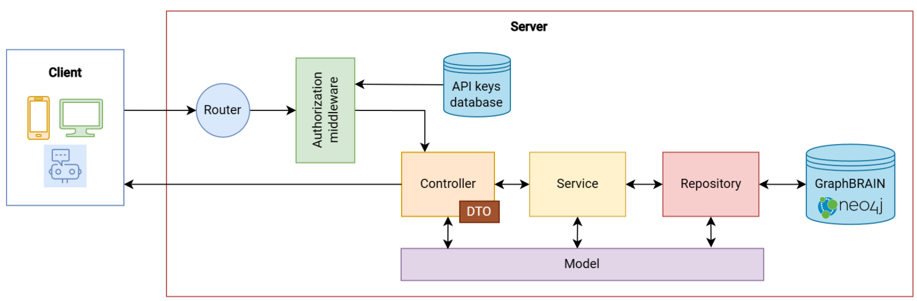

Fig. 3. NutriWell API implementation architecture

The framework is deployed leveraging containerization with two Docker containers (one for the API server, and the other for the Neo4j database) based on custom-built Docker images.

6 Expert Evaluation and Use Case

The evaluation has been conducted with the help of two members of the Department of Science, Food and Nature of the university. They have been chosen of almost the same age and different genders. They were responsible for identifying and pointing out possible misleading values in the labelling phase (Section 3.2) and adding further considerations. They were allowed to modify the impact of some aliments on some diseases. Throughout the dataset, the main rationale was that the aliment is the sum of all its ingredients, but this is not always the case. Unfortunately (or luckily) recipes are not standard and this is the first element hindering the accuracy of the solution. For this reason, experts tended to be less permissive when considering some aliments, because it happens that the recipes introduce relationships among ingredients that do not exist when taken individually. The two experts were separated during their evaluation of the labels. At the end of their individual examinations, they exchanged their results and discussed them to reach an agreement. In the end, less than 2% of the impact changed polarity (from 0 to -0.5 or -1). In no case did the polarity change from a positive value to a negative one or vice-versa. One of the evaluators made very few changes from a negative to a neutral impact, but these changes have not been maintained after discussion with the other evaluator. The first evaluator proposed changes in 0.8% of the dataset, while the second in the 1.1%. After the discussion, 1.5% circa was changed. The agreement between the two evaluators after individual analysis was about 40%, which is quite low but not unexpected. The difficulty in the agreement lies in the subjectivity of the evaluation. Factors like quantity, ingredient matching and personal experience come into play. Nonetheless, an agreement was smoothly found and the small percentage of values changed justifies the use of common knowledge to label data.

6.1 Use Case

We present here a stereotypical situation in which an end-user may benefit from the system. Figure 4 reports the interaction with the application querying *NutriWell*.

Scenario *Mario is a retired 75-year-old man living in Italy and, despite being widowed and living alone, he leads a quite active life. However, his fondness for sweets has led to health issues in recent years. Indeed, he has been diagnosed with irritable bowel syndrome and high cholesterol. Concerned for his health, he installed on his smartphone a mobile app with diet-tracking capabilities to assist during mealtimes where he can report his health conditions.*

When Mario asks the system to check whether a box of chocolates may fit with its health, he observes that chocolates hurt irritable bowel syndrome (-0.7*), but do not affect high cholesterol (*0.0*). The app also explains to Mario which are the harmful ingredients in chocolates. Thus, he decides to eat only one.*

Then, he asks for suggestions on more appropriate meals for his condition. The system looks for meals having at least an impact score equal to 0.5 *and presents them to Mario.*

7 Ethical Concerns

Although the scraping of *GialloZafferano* and *AlimentiNUTrizione* is not explicitly allowed, we want to remind you that *NutriWell* is not a recipe framework, it

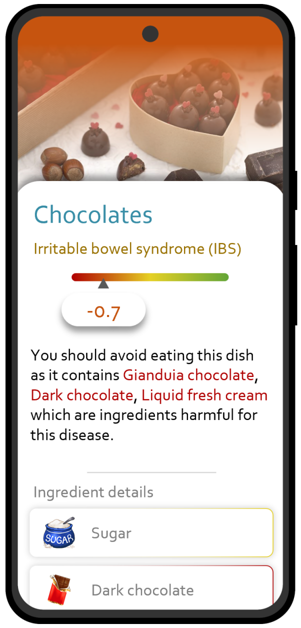

Fig. 4. Example of *NutriWell* response for chocolates

uses information about recipes to categorise meals and expose potential risks for users. Under no circumstances could the user get where the information about potential harm comes from, and hence information about recipe authors can never be extracted, since we discard it in the dataset pre-processing.

8 Conclusions

In this study, we proposed a novel approach to assist end-users in choosing foods for both preference or health motivations, using and combining structural components of food and ingredients. We provided the assistance in the form of a REST API and contributed to the definition of a new conceptualization for health conditions management. The solution employs the *GraphBRAIN* framework. This work may give rise to several extensions, from the ontological formalization of inter-relationships among ingredients to the definition of exceptions when dealing with aliments' interactions, which is not uncommon. Data labelling may require further analysis in the future but, as far as the two experts were concerned, no relevant criticalities emerged. In addition, a thorough evaluation of the application prototype with end users will be conducted to understand its impact on improving nutrition literacy and awareness.

9 Citations and references

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Angles, R.: The Property Graph Database Model. In: Alberto Mendelzon Workshop on Foundations of Data Management (2018), https://api.semanticscholar.org/CorpusID:43977243
2. Bajaj, V., Panda, R.B., Dabas, C., Kaur, P.: Graph Database for Recipe Recommendations. In: 2018 7th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO). pp. 1–6 (2018). https://doi.org/10.1109/ICRITO.2018.8748827
3. Çelik, D., et al.: Foodwiki: Ontology-driven mobile safe food consumption system. The scientific World journal **2015** (2015)
4. Cena, H., Calder, P.C.: Defining a healthy diet: evidence for the role of contemporary dietary patterns in health and disease. Nutrients **12**(2), 334 (2020)
5. Dragoni, M., Bailoni, T., Maimone, R., Eccher, C.: Helis: An ontology for supporting healthy lifestyles. In: The Semantic Web–ISWC 2018: 17th International Semantic Web Conference, Monterey, CA, USA, October 8–12, 2018, Proceedings, Part II 17. pp. 53–69. Springer (2018)
6. Edaman: Edaman API (2024), https://www.edamam.com/, [Accessed: 17-Apr-2024]
7. Federation, I.D.: IDF Diabetes Atlas (2024), https://diabetesatlas.org/data/en/region/3/eur.html, [Accessed: 17-Apr-2024]
8. Federation, I.D.: IDF Diabetes Atlas (2024), https://www.who.int/news-room/factsheets/detail/noncommunicable-diseases, [Accessed: 17-Apr-2024]
9. Ferilli, S., Redavid, D.: The GraphBRAIN system for knowledge graph management and advanced fruition. In: Foundations of Intelligent Systems: 25th International Symposium, ISMIS 2020, Graz, Austria, September 23–25, 2020, Proceedings. pp. 308–317. Springer (2020)
10. Ferilli, S., Redavid, D., Di Pierro, D., et al.: LPG-based Ontologies as Schemas for Graph DBs. In: SEBD. pp. 256–267 (2022)
11. Harrington, J.L.: Relational database design and implementation. Morgan Kaufmann (2016)
12. Kamel Boulos, M.N., Yassine, A., Shirmohammadi, S., Namahoot, C.S., Brückner, M.: Towards an "Internet of Food": food ontologies for the internet of things. Future Internet **7**(4), 372–392 (2015)
13. Mavani, N.R., Ali, J.M., Othman, S., Hussain, M., Hashim, H., Rahman, N.A.: Application of artificial intelligence in food industry—a guideline. Food Engineering Reviews **14**(1), 134–175 (2022)
14. Popovski, G., Korousic-Seljak, B., Eftimov, T.: FoodOntoMap: Linking food concepts across different food ontologies. In: KEOD. pp. 195–202 (2019)
15. Qin, L., Hao, Z., Zhao, L.: Food safety knowledge graph and question answering system. In: Proceedings of the 2019 7th International Conference on Information Technology: IoT and Smart City. pp. 559–564 (2019)
16. Rakhmawati, N.A., Fatawi, J., Najib, A.C., Firmansyah, A.A.: Linked open data for halal food products. Journal of King Saud University-Computer and Information Sciences **33**(6), 728–739 (2021)

17. Sagan, A., Kowalska-Bobko, I., Bryndová, L., Smatana, M., Chaklosh, I., Gaál, P.: What is being done to respond to the rise of chronic diseases and multi-morbidity in Czechia, Hungary, Poland, and Slovakia? Frontiers in Public Health **10**, 1082164 (2023)
18. Silva, P.: Food and nutrition literacy: Exploring the divide between research and practice. Foods **12**(14), 2751 (2023)
19. Snae, C., Bruckner, M.: FOODS: a food-oriented ontology-driven system. In: 2008 2nd ieee international conference on digital ecosystems and technologies. pp. 168–176. IEEE (2008)
20. Toledo, R.Y., Alzahrani, A.A., Martinez, L.: A food recommender system considering nutritional information and user preferences. IEEE Access **7**, 96695–96711 (2019)
21. Trattner, C., Elsweiler, D.: Food recommender systems: important contributions, challenges and future research directions. arXiv preprint arXiv:1711.02760 (2017)
22. Vlachou, A., Stavroussi, P., Roka, O., Vasilou, E., Papadimitriou, D., Scaratti, C., Kadyrbaeva, A., Fheodoroff, K., Brecelj, V., Svestkova, O., et al.: Policy guidelines for effective inclusion and reintegration of people with chronic diseases in the workplace: national and european perspectives. International Journal of Environmental Research and Public Health **15**(3), 493 (2018)
23. World Health Organization: Noncommunicable diseases (2024), https://health.ec.europa.eu/non-communicable-diseases/overview_en, [Accessed: 17-Apr-2024]
24. Yu, E., Rimm, E., Qi, L., Rexrode, K., Albert, C.M., Sun, Q., Willett, W.C., Hu, F.B., Manson, J.E.: Diet, lifestyle, biomarkers, genetic factors, and risk of cardiovascular disease in the nurses' health studies. American journal of public health **106**(9), 1616–1623 (2016)