Assessing the Validity of a Functional Status Knowledge Graph in a Large-Scale Living Lab

Mauro Dragoni[1,*], Gianluca Apriceno[1,*] and Tania Bailoni[1,*]

[1]*Fondazione Bruno Kessler, Via Sommarive 18, Povo, 38213, Italy*

Abstract

Functional Status Information describes physical and mental wellness at the whole-person level. Collecting and analyzing this information is critical to address the needs for caring for an aging global population, and to provide effective care for individuals with chronic conditions, multi-morbidity, and disability. Knowledge Graphs represent a suitable way for meaning in a complete and structured way all information related to people's Functional Status Information and reasoning over them to build tailored coaching solutions supporting them in daily life for conducting a healthy living. In this paper, we describe the integration of our Functional Status Knowledge Graph, namely FuS-KG, into a real-world application run within a large-scale living lab involving more than 4,000 people. We provide the road map of this experience including the challenges, the platform's architecture, the focus on the knowledge layer, and the evaluation and insights observed.

Keywords

Knowledge Graph, Digital Health, Large-scale Living Lab

1. Introduction

There is a growing trend of developing virtual health and well-being assistants to support lifestyle and disease management, partly due to the growing societal need for managing health and preventing illness. To improve an individual's situation, a change of behavior is typically necessary, which puts focus on how a digital coach can act in collaboration with the individual to support the individual's ambition to improve their health through behavior change, e.g., by adhering to medical guidelines or treatment protocols, increasing physical activity, changing nutrition habits, reducing stress or intake of toxic substances. As a basis for deciding how to act, the digital coach may explore the Functional Status Information (FSI) of monitored individuals.

A necessary foundation for a medical and health-related system's reasoning, decision-making, and acting is (i) the medical knowledge that the digital coach utilizes, (ii) the theories and knowledge about how humans form motivation and change behavior as well as manage physical, social, and psychological barriers, and (iii) the FSI (data) about the individual as well as the individuals' narrative about their behavior change journey, information that needs to be treated following ethical guidelines and regulations. Moreover, the realization of such systems relies on the integration of effective, efficient, and ethical strategies for adapting behavior in a situation depending on the individual's context, personal preferences, and needs (e.g., displaying motivational messages that are tailored to each individual's resources and current situation).

A proper representation of this information requires a strategy able to mitigate the diversity of the information managed and a conceptual model enabling the exploitation of such information by preserving, at the same time, the privacy aspects. Knowledge Graphs (KGs) are a valid way to provide an effective representation of FSI and to connect such information with users' records (e.g., electronic health records) to enable the design of AI-based systems implementing the coaching paradigm for avoiding Functional Status (FS) deterioration in target users.

HC@AIxIA 2024: 3rd AIxIA Workshop on Artificial Intelligence For Healthcare, November 25–28, 2024, Bolzano, Italy
*Corresponding author.
dragoni@fbk.eu (M. Dragoni); apriceno@fbk.eu (G. Apriceno); tbailoni@fbk.eu (T. Bailoni)
0000-0003-0380-6571 (M. Dragoni); 0000-0003-4603-6888 (G. Apriceno); 0000-0003-4360-566X (T. Bailoni)

In this paper, we present the experience concerning the adoption into a real-world setting of a brand new KG, namely FuS-KG, supporting the representation of FSI and its exploitation to manage the generation of healthy recommendations for citizens.

Our experience is split into the three major steps presented in the remainder of the paper and summarized below. First, the construction of FuS-KG (Section 3), which represents the first original contribution of this work. Second, the integration of FuS-KG into a real-world digital health platform (Section 4), namely Salute+, is provided for completeness of understanding, but it must not be considered an original contribution of this paper. Third, the validation of the FuS-KG-enhanced Salute+ platform within a large-scale living lab involving more than 4,000 users within the Trentino territory (Section 6), which is the second original contribution of this work demonstrating the effectiveness of the proposed solution.

2. Related Work

Understanding the FS of a person is a key step toward tailoring precise interventions and providing services to improve an individual's health status and maximize functional independence in daily activities. However, managing health often poses different challenges, related to healthcare systems being unable to provide adequate support to individuals with chronic illnesses or disabilities due to the scarcity of resources. The National Committee on Vital and Health Statistics (NVHCS) states that understanding the FS is necessary to achieve optimal health outcomes [1]. The evaluation of the FS is characterized by a clinician-patient interaction to get insights into patients' lifestyles, together with standardized tests. Nevertheless, these assessments provide merely a snapshot of an individual's health status, while effective behavior change for health improvements is a long-term process. Consequently, a gap exists between an individual's health objectives and the ''infrequent'' health assessments conducted in healthcare settings, which may prevent the delivery of high-quality care. Therefore, frequent data collection is required to bridge this gap and ensure continual support.

Physicians and researchers have conducted a few studies on elderly people that led to the identification of the risk factors able to detect those most susceptible to functional decline [1, 2, 3]. Nonetheless, FSI is underutilized, limiting its potential impact [4].

One of the reasons may be that many physicians do not fully recognize the significance of this information [5, 6, 7, 8]. Indeed, even when informed about a patient's health status, only a minority alter patient management accordingly.

Studies based on self-reports of functional performance and early decline show successful prediction of actual performance and decline [9, 10]. However, researchers note these reports capture only a fraction of the problems [3]. Consequently, there is a need for innovative methods to assess FS, including non-intrusive approaches that evaluate the patient's functional performance, in both normal and abnormal conditions, without relying solely on physician-patient interaction. For example, in [11] sensor-based systems can detect a decline in daily activities, thus facilitating tailored interventions. Given the increasing elderly population, such systems are crucial for early detection and intervention to prevent functional deterioration and reduce the decline of functional ability.

FuS-KG enables the investigation of this research field aiming to facilitate the development of AI-powered systems capable of helping individuals monitor their FS and prevent function decline. The adoption of FuS-KG in real-world scenarios is strongly connected to the design of behavior change solutions since, given the detection of declining FS, the knowledge modeled within FuS-KG can be exploited to perform reasoning operations for discovering possible behavior change trajectory with the aims of improving the overall FS of a person. Concerning this point, FuS-KG fills a gap in the state-of-the-art since, as described below, to the best of our knowledge, KGs covering the whole FS landscape have not been proposed yet.

An example of a recent collaborative effort to develop observable and replicable interventions for influencing behavior and health can be found in the behavioral change techniques taxonomy outlined by [12]. This taxonomy facilitates the consolidation of knowledge regarding behavior and behavior change,

[1]https://www.ncvhs.hhs.gov/wp-content/uploads/2017/08/010617rp.pdf

promoting the sharing and reuse of useful sources of behavior knowledge. Additionally, we note that, as far as we are aware, at present there is no publicly accessible conceptual classification that conceptualizes barriers to behavior change.

An important taxonomy to mention is the human behavior taxonomy developed by the World Health Organization [2] (WHO) [13]. This taxonomy, rooted in the WHO's expertise and the International Classification of Functioning (ICF), Disability, and Health, provides detailed class definitions based heavily on the U.S. National Cancer Institute (NCI) Thesaurus, as well as the Oxford English Dictionary [14]. Another important effort aimed at understanding human behavior is the Semantic Mining of Activity, Social and Health Data (SMASH) [15]. This project focuses on predicting human behavior and providing explanations for these predictions.

The Health Behavior Change Ontology (HBCO) was developed for a project to establish automated dialogues between a psychologist and a user to provide behavioral counseling [16]. While the HBCO ontology has enhanced the connection between theoretical and practical aspects, there are limited practical implementations available. In turn, there is currently a lack of specific strategies for deploying a reusable behavior change ontology in practice.

Since our conceptualization has to align with the individual undergoing behavioral change, it is fundamental to explore ontologies specifically modeling users, encompassing their profiles, characteristics, and sometimes their behavior. For example, the General User Model Ontology (GUMO) [17] (included in our ontology), the User Navigation Ontology (UNO) [18] and the Ontology of Personal Information Management (OntoPIM) [19].

Moreover, to cover the domain of physical activity behavior, notable examples of behavior ontologies related to physical activity or exercise are the Ontology of Physical Exercises [3] and the HeLiS ontology [20]

Finally, the use of KGs in the design of AI-based systems providing coaching has been researched and evaluated. [21] describes a knowledge-based solution implemented in a system providing personalized healthy lifestyle recommendations to users which was applied and evaluated in a real-world scenario. The article highlights the already mentioned lack of structured knowledge publicly available but also provides evidence of the feasibility of the proposed solution in a real-world setting based both on performance evaluation and efficacy.

3. Building FuS-KG

The development of FuS-KG addressed a list of requirements aiming to satisfy the need to provide a KG able to describe the domain of FS as-is and to support the conceptualization of an interoperable digital twin of people using applications exploiting FuS-KG. This way, FuS-KG enables the possibility to perform reasoning operations on such a digital twin also for predictive purposes by simulating possible future undesired behaviors and, in turn, by generating recommendations persuading the users to avoid them.

As we discussed in Section 2, ontologies available in domains connected with wellness and healthy lifestyle have been designed with different aims. For example, ontologies concerning the physical activity domain are created with a focus on classifying data without connecting each activity with potential problems or benefits associated with the whole FS of a user. Similarly, the same happened with ontologies concerning foods. FuS-KG has been built with a focus on the connection between different dimensions related to people's health representing their complete FS. The development of our ontology has been driven by the following requirements:

Requirement 1 (REQ1). The KG must conceptualize the food domain at a fine-grained level. This means that the whole knowledge chain from defining each nutrient to modeling complex recipes must be supported.

Requirement 2 (REQ2). The KG must provide a comprehensive list of activities that a person can perform. For each activity, it must be provided the knowledge required to understand the effort necessary

[2]https://www.who.int/classifications/drafticfpracticalmanual.pdf
[3]https://bioportal.bioontology.org/ontologies/OPE/?p=summary

to complete the activity to enable the inference of how much of each food defined under *REQ1* is necessary to fulfill the activity.

Requirement 3 (REQ3). The KG must include the model of the barriers that may affect a person and how such barriers obstacle the fulfillment of specific activities, the consumption of specific foods, or, in general, the following of specific guidelines.

Requirement 4 (REQ4). The KG must support the modeling of multi-modal knowledge since a MMKG may be exploited to better support users' education tasks and enable knowledge injection tasks into large foundational models. Hence, the KG must include a multi-modal knowledge representation, e.g., images of recipes and videos of how to execute activities.

Requirement 5 (REQ5). Knowledge modeled under the requirements *REQ1*, *REQ2*, *REQ3*, and *REQ4* must be associated with knowledge and data gathered from users' input. Hence, the KG must include the appropriate set of concepts to enable the definition of a user model and to support the linking between such a user model and the domain knowledge defined through the requirements mentioned above.

Requirement 6 (REQ6). A KG usable for creating a behavior change solution requires a set of guidelines driving the behavior change intervention. A reasoner can exploit such guidelines to detect situations where a user is not adhering to them. To enable this feature the KG must define the conceptual knowledge that: (i) enables the modeling of such guidelines; (ii) defines how they can be associated with the domain knowledge covered by the KG; and, (iii) allows their linking with a user profile.

Requirement 7 (REQ7). The requirements discussed above refer to the notion of *time* in different ways. For example, the merge of *REQ2* and *REQ4* concerning the modeling of an activity and how to perform it, requires the modeling of the different steps and their temporal order. Similarly, both requirements *REQ5* and *REQ6* need the notion of time associated with the knowledge gathered by the user and when users' data should be checked, respectively. Hence, the KG must include temporal knowledge to support the requirements above and enable temporal reasoning over the users' collected knowledge.

Requirement 8 (REQ8). Working with such different domains may lead to building a very large KG. Hence, the KG must be designed with a modular structure to ease its management and maintenance.

The process for building FuS-KG followed a combination of the Modular Ontology Modeling (MOMo) [22] and the METHONTOLOGY [23] methodologies. The rationale behind applying them in tandem is that in the first phase, we worked on the modularization aspect given the expected size of FuS-KG and the purpose of easing the possible reuse of only parts of the knowledge modeled. Then, we moved to the conceptualization of each module. The choice of METHONTOLOGY was driven by the necessity of adopting a lifecycle split into well-defined steps. Moreover, the development of FuS-KG requires the involvement of the experts in situ. Thus, the adoption of a methodology having a clear definition of the tasks to perform was preferred. Other methodologies, like DILIGENT [24] and NeOn [25], were considered before starting the construction of FuS-KG. However, the characteristics of such methodologies, like the emphasis on decentralized engineering, did not fit well our scenario.

The overall process involved four knowledge engineers and three domain experts from the Trentino Healthcare Department. More precisely, three knowledge engineers and two domain experts participated in the ontology modeling stages (hereafter, the modeling team). While, the remaining knowledge engineer and domain expert were in charge of evaluating the ontology (hereafter, the evaluators).

Due to limited space, we do not provide a detailed description of each phase, but we limit ourselves to reporting the most relevant activities of the construction process.

3.1. Modules Definition

As mentioned above, the first step focused on the application of the MoMo methodology to define FuS-KG modules. This step aims to address *REQ8*. This step was necessary since the amount of knowledge created by starting from the selected unstructured sources was huge. Hence, to ease the maintenance of FuS-KG, the split into a set of modules was a mandatory step. Figure 1 provides an overview of the modules composing the current version of FuS-KG.

The knowledge contained in each module is the following:

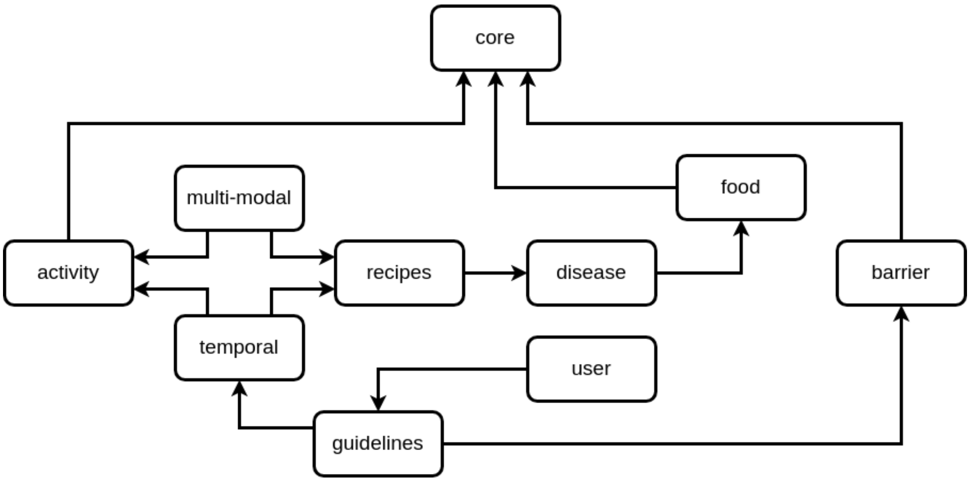

Figure 1: Summary of the modules composing the current version of FuS-KG. Each block represents a module of FuS-KG, while each arrow corresponds to the *owl:imports* property.

- *core*: this module includes the upper level of FuS-KG, i.e., the set of abstract concepts defining the main type of knowledge covered by FuS-KG.
- *food*: this modules imports the *core* module and it contains all the knowledge about the *BasicFood*, *ComposedFood*, and *Nutrient*. As *BasicFood*, we mean those foods for which fine-grained nutritional information is provided within the sources we adopted to build FuS-KG (e.g., Bread). Then, as *ComposedFood*, we mean those foods that are aggregations of instances of *BasicFood* but for which fine-grained nutritional information is available as well (e.s., Tomato Sauce). While *Nutrient* represents the specific nutritional information associated with a *BasicFood*.
- *recipes*: this module imports the *diseases* and (indirectly) the *food* module since a *Recipe* is defined as a group of *BasicFood* each with the associated quantity. This module includes a set of sub-modules containing (i) the list of the recipes collected from the different sources (we created one sub-module for each source); and, (ii) the alignments between the different sources (i.e., we preserve the fact that a recipe may be defined within more than one of the sources we used).
- *diseases*: this module imports the *food* module and it models the association between each *BasicFood* and nutritional-wise *Disease*. Each association is modeled through a *DiseaseRiskLevel* entity associating to each pair *BasicFood-Disease* a risk level.
- *activity*: this module imports the *core* module and it includes the taxonomy of all activities covered by FuS-KG together with the knowledge related to the effort required to fulfill each activity. The effort is represented through the Metabolic Equivalent of Task (MET) coefficient and the calories required to perform 1 minute of each activity for each kilogram of body weight.
- *barrier*: this module imports the *core* module and it contains the type of barriers defined within the SIS manual and how these barriers may affect the fulfillment of specific activities.
- *multi-modal*: this module imports the *recipes* and *activity* modules. The module contains, when available within the adopted sources, the multi-modal knowledge associated with specific *Recipe* and *Activity*. In particular, the current version of FuS-KG contains knowledge about image and video modalities.
- *temporal*: this module imports the *recipes* and *activity* modules. Then, this module is split into two sub-modules. The first one defines the temporal intervals that may be of interest to model the guidelines (e.g., *Day*, *Meal*). While the second one contains the knowledge related to steps to prepare a specific *Recipe* or perform a specific *Activity*.
- *guidelines*: this module imports the *barrier* and *temporal* modules. This module aims to model behavioral guidelines that may be associated with users to support reasoning tasks about possible recommendations related to the user data that can be collected [26].
- *user*: this module imports the *guidelines* module. This way, the *user* module may access the full

knowledge of FuS-KG to enable the storage of any type of data that can be collected in several ways (e.g., sensors or mobile applications).

3.2. Modules Conceptualization

FuS-KG aims to provide a conceptualization with a *high* granularity level. For example, for each recipe modeled within FuS-KG, we provided its composition to the micro-nutrient level. Thanks to this granularity level, we favor the integration of FuS-KG into several solutions going from simple mobile applications to expert systems.

The acquisition of the knowledge necessary for building FuS-KG has been done in two steps: (i) the discussion with domain experts for deciding how to model the core entities of FuS-KG (i.e., abstract classes and properties); and, (ii) the acquisition, analysis, and processing of unstructured resources containing information to include in FuS-KG.

The first step consisted of defining the set of entities addressing the list of requirements described above. Here, the modeling team started from the conceptual model built to create the HeLiS ontology [27] since it provides the basic elements to start satisfying requirements *REQ1*, *REQ2*, *REQ5*, and *REQ6*. Then, such a conceptual model has been extended with further entities defined by the modeling team to cover the remaining requirements, i.e., *REQ3*, *REQ4*, and *REQ7*. This way, the final description of the core FuS-KG conceptual model defines the barrier domain and it has been equipped with the capabilities of accepting knowledge related to both multi-modal resources and temporal information about entities.

The second step consisted of identifying the sources for building and populating FuS-KG. In particular, such sources were related to the following domains: food, activity, and barriers. The HeLiS ontology has been used as a starting point to build the T-Box of FuS-KG. Concerning the food domain, we imported into the FuS-KG schema, the model of the recipes already defined within the HeLiS ontology, i.e., (i) the archives of the Italian Minister of Agriculture [4] and the Italian Epidemiological department [5]; and, (ii) the Turconi's atlas [28]. Then, we selected the following four (1 structured and 3 unstructured) further sources concerning the food and nutrition domain: (i) the USDA database [6], enriching the lists of basic foods and nutrients provided by the HeLiS ontology; (ii) the Recipe1M [7] dataset, which provides both images and step-wise description of recipes; (iii) the Tasty [8] dataset, which provides descriptions and videos of recipe preparations; and, (iv) the RecipeDb [9] dataset, which provides a comprehensive set of recipes still missing in FuS-KG. During the acquisition of such information, we also worked on the alignment between the INRAN and USDA models to have a common representation of basic foods and nutrients. This way, it was possible to reconcile the different sources exploited to build FuS-KG.

Concerning the activity domain, we started from the Compendium of Physical Activities [10] to create the taxonomy of physical activities and model all information concerning the associated effort. Finally, concerning the barrier domain, we relied on the Supported Intensity Scale (SIS) manual [11] that provided all the knowledge necessary to model barriers used to measure the functional status of a person. Here, we integrated object properties defining which barriers may affect the capability of fulfilling specific activities.

As a final step, we focused on the refinement of the conceptual model adopted within FuS-KG and the definition of the ontology design patterns (ODP) [29] to adopt. From the ODP catalog [12], we adopted several patterns, in particular: the logical patterns *Tree* and *N-Ary Relation*, the alignment pattern *Class Equivalence*, and the content patterns *Parameter*, *Time Interval*, *Action*, *Classification*.

[4] http://nut.entecra.it/
[5] http://www.bda-ieo.it/
[6] https://fdc.nal.usda.gov/
[7] http://pic2recipe.csail.mit.edu/
[8] https://cvml.comp.nus.edu.sg/tasty/
[9] https://cosylab.iiitd.edu.in/recipedb/
[10] https://pacompendium.com/
[11] https://www.aaidd.org/sis
[12] http://ontologydesignpatterns.org/wiki/Community:ListPatterns

Finally, the main metrics related to the content of FuS-KG are summarized to give an overview of its size: 588 Concepts, 128 Object Properties, 49 Data Properties, 58 Annotation Properties, 1879205 Individuals, and 8668859 Axioms.

4. Integration Into the Salute+ Platform

FuS-KG has been integrated into the Salute+ platform to enable its usage. In this section, we provide a brief description of how FuS-KG has been integrated and the role it plays. This part is reported for completeness of the in-use experience described in this paper, but it is not an original contribution of this work. The reader may refer to [26] for a more in-depth description of the general architecture of the platform and about the message generation pipeline.

The Salute+ platform is composed of the four layers shown in Figure 2. The *Input Layer*, is responsible for storing events that trigger the platform activities and accounts for the system's ability to sense the context of interaction. These events are of two types: (i) data input, where data are sent from the *Input Layer* to the *Knowledge Layer*, and (ii) context communication, where contextual information is sent from the *Input Layer* to the *Communication Layer* that may exploit this information for communication purposes. The *Knowledge Layer* encompasses the FuS-KG resource described in Section 3. The *Communication Layer* exploits the output of the *Knowledge Layer* (i.e., reasoning operations) for choosing the language strategies to include in the natural language-generated messages, and focuses on the tasks of selecting the arguments to include in the message, to order them, and to choose the right wording for each argument. More information is provided below, but a detailed description of this layer is out of the scope of this paper. Finally, the *Output Layer* closes the loop by providing the generated message to users. It represents the many devices that can receive the data produced by the *Communication Layer* and conveys the physical feedback to users.

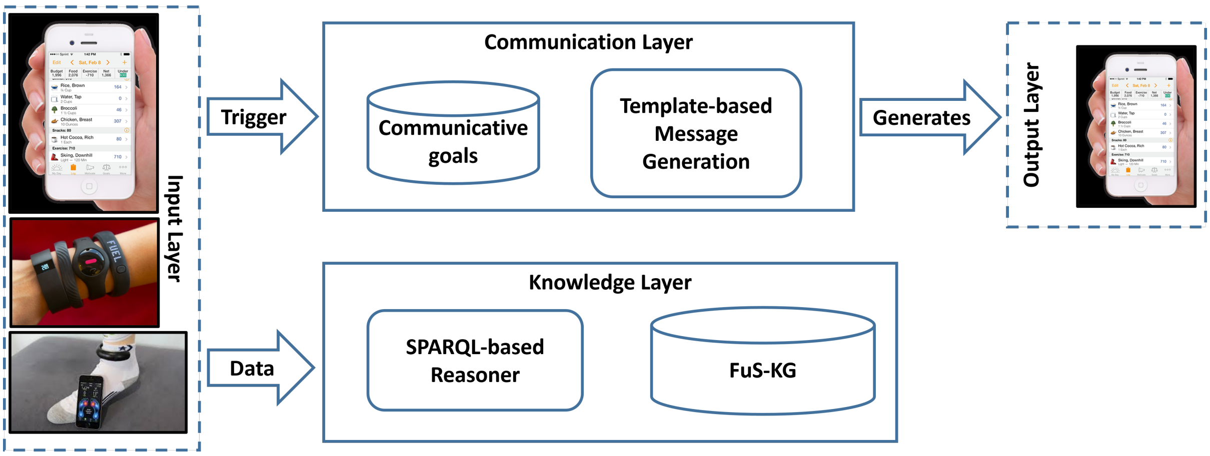

Figure 2: Overview of the Salute+ architecture.

FuS-KG is exploited for monitoring the functional status of a user through its integration into a SPARQL-based reasoner. Such a reasoner is used for detecting undesired situations within users' behaviors. When inconsistencies between the data provided by a user and the associated guidelines are detected, the reasoner generates an individual of type *UndesiredEvent* that is then exploited by the *CommunicationLayer* to generate feedback to users. The reasoner activity can be triggered in two ways. Firstly, each time a new, or updated, data package is provided by a user (or acquired by an IoT device), the reasoner processes the new, or updated, information. Secondly, at the end of a specific timespan (e.g., the end of a day, or the end of a week), the reasoner checks data concerning the behavior of each user included in the system in such a timespan. In the latter case, the reasoner works on a collection of data labeled with a timestamp valid within the considered timespan. The integrated reasoner relies on the architecture implemented in RDFpro [30]. RDFpro has been chosen for two main reasons. Firstly, the architecture of RDFpro allows the integration of custom methods into reasoning operations (i) for performing mathematical calculations on users' data and (ii) for exploiting real-time information acquired from external sources without materializing them within the knowledge repository. Secondly, as reported in [30], efficient analysis performed on RDFpro demonstrated the suitability of this reasoner compared with other state-of-the-art reasoners in a real-time scenario. In this work, RDFpro has been adapted and extended to better fit the needs of the proposed solution. The extension consisted of the integration of new methods supporting the real-time stream reasoning of sensor data. This way, we were able to support the real-time processing of users' data more efficiently.

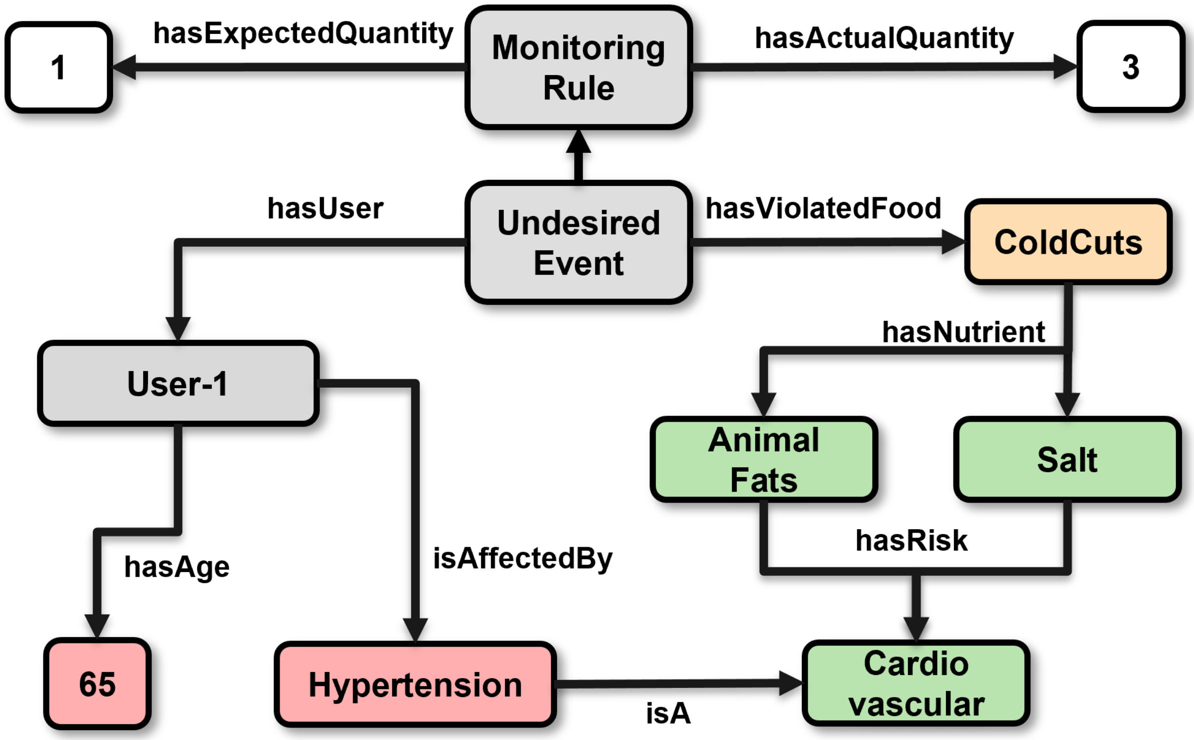

Figure 3: Example of an *UndesiredEvent* individual generated by the reasoner.

Figure 3 shows an example of the knowledge that is generated when an undesired event is detected. A new individual of type *UndesiredEvent* is created and it is linked to the guideline that has been violated, the data led to the undesired event, and data coming from the user profile.

The generated knowledge is given as input to the *Communication Layer* that applies a template-based strategy to generate the recommendation sent to the user. The example shown in Figure 3 is translated into the following message: ''This week you consumed too much (3 portions of a maximum 1) cold cuts. Cold cuts contain animal fats and salt that can cause cardiovascular diseases. People over 60 years old are particularly at risk especially if they suffer from hypertension. Next time try with some fresh fish''.

The large-scale living lab described in Section 6 aimed to demonstrate the effectiveness of the messages generated by using FuS-KG compared with pre-defined canned texts and the messages generated by using the HeLiS ontology (i.e., the ontology we presented in [27] containing only knowledge about food and physical activities).

5. Large-scale Living Lab: The Salute+ Project

As introduced in Section 1, this work is part of the Salute+ [13] project. The project is a set of innovative interventions/initiatives aimed at offering the general population and specific categories of subjects,

[13]https://trentinosalutedigitale.com/blog/portfolio/trentinosalute/

different functions to create a context favorable to healthy choices through concrete opportunities and tools to counteract the risk factors for the onset of chronic diseases such as poor nutrition and a sedentary lifestyle. This project is part of a set of citizen health promotion initiatives run by *Trentino Salute 4.0* [14]. Within this framework, we ran a territorial living lab that enrolled 4,274 citizens using the mobile application integrated within the Salute+ platform for seven weeks.

The enrollment of the users was voluntary and no incentives were adopted. Indeed, all the users were already motivated to participate in the study. All involved users have been equipped with smart bands that synchronize information about steps and physical activity data with our system. However, users were asked to insert a report of performed activities also manually to validate the synchronized information. Part of future work is to reduce the effort in the acquisition of physical activity information.

This study represented the first large-scale living lab in Trentino concerning the adoption of knowledge-equipped AI tools supporting health prevention. The *Trentino Salute 4.0* aims at extending the adoption of the Salute+ solution to at least 30,000 citizens of the Province of Trento before the end of 2024 and to open the adoption of this platform at the Italian National level during the two years 2025-2026.

For completeness, Table 1 shows the main demographic information concerning the citizens involved in the evaluation proposed in this work. We want to highlight that in this living lab, all users presented a healthy status since in this first pilot, we decided not to accept people affected by chronic or other diseases.

Dimension	Property	Value
Gender	Male	57%
	Female	43%
Age	25-35	12%
	36-45	49%
	46-55	23%
	56-65	16%
Education	High-school or lower	42%
	University Degree	58%
Type of Occupation	Sedentary	39%
	Active Indoor	34%
	Active Outdoor	27%

Table 1
Distribution of demographic information of the users involved in the evaluation campaign.

6. Evaluation

In this Section, we report the evaluation activities we performed on the Salute+ solution during the forty-nine days timespan of the large-scale living lab described in Section 5. The results observed from the collected data are presented in Section 6.1. While Section 6.2 discusses the key lessons learned from this experience.

6.1. Results

This user study consisted of providing a group of users with a mobile application we created based on the services included in the Salute+ solution. In particular, we measure the effectiveness of the recommendations generated through the use of FuS-KG and we compared them with the results we reported in [26] where a non-randomized experiments setup [31] was conducted with a Control Group received predefined canned text messages and an Intervention Group received messages created with the same methodology adopted in the Salute+ by relying on the HeLiS ontology instead of FuS-KG.

The sets of guidelines implemented in this living lab were the same as used in the baselines allowing a fair comparison of the results:

[14] https://trentinosalutedigitale.com/

- MEAL-Rules (related to single meals) that check the correct quantity of a specific food category to be consumed in a single meal. Users were asked to insert 4 meals every day: breakfast, lunch, snack, and dinner.
- DAY-Rules (related to a single day) that check the maximum (or minimum) quantity (or portion) of a specific food category that can (or should) be daily consumed.
- WEEK-Rules (related to a single week) that check the maximum (or minimum) quantity (or portion) of a specific food category that can (or should) be weekly consumed.

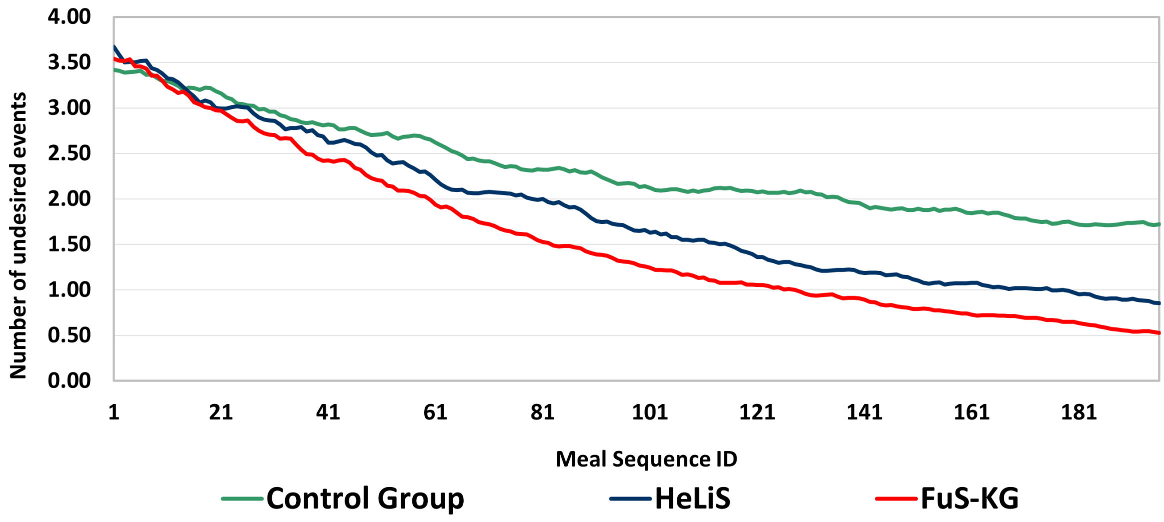

Figure 4: Trend of the number of undesired events observed on the MEAL-Rules.

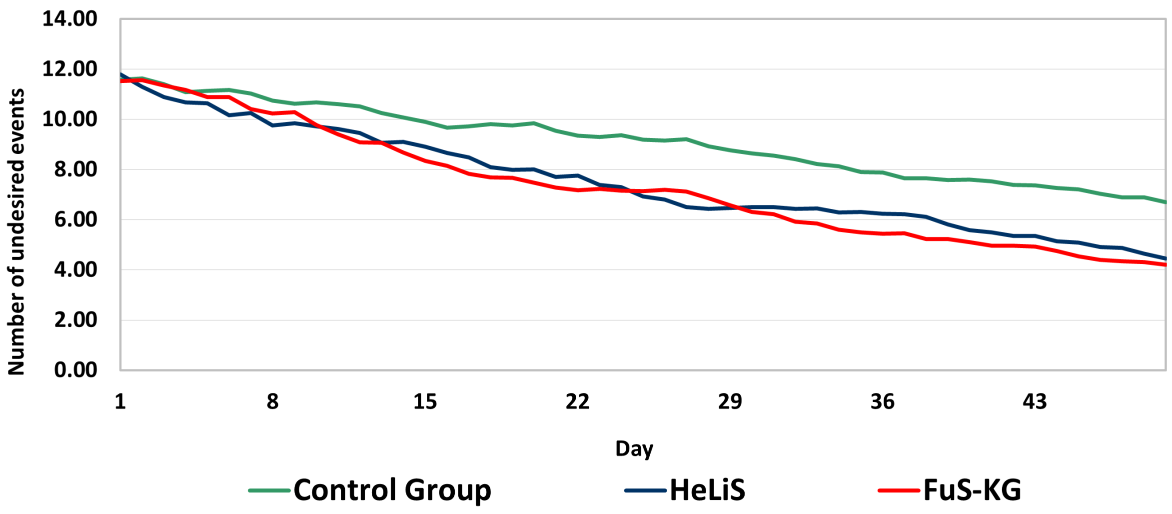

Figure 5: Trend of the number of undesired events observed on the DAY-Rules.

Figures 4, 5, and 6 present the evolution of the average number of undesired events per user detected concerning the MEAL-Rules, DAY-Rules, and WEEK-Rules sets, respectively. The green lines represent the trends of the Control Group, the blue lines represent the trends of the Intervention Group received recommendations generated by using the HeLiS ontology, and, finally, the red lines represent the trends observed on the users who adopted the Salute+ solution.

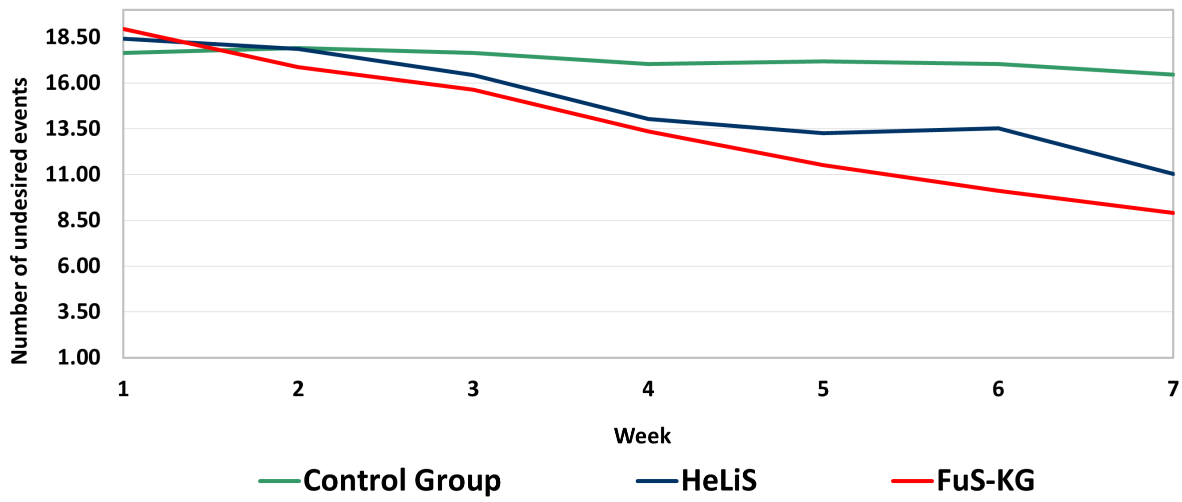

Figure 6: Trend of the number of undesired events observed on the WEEK-Rules.

As mentioned above, MEAL-Rules are verified every time a user enters a meal within the system; DAY-Rules are verified at the end of the day; while WEEK-Rules are verified at the end of each week. The increasing gap between the green lines and the others demonstrates the positive impact of the knowledge-based generated recommendations sent to users. We can observe how the average number of detected undesired events for the MEAL-Rules is below 1.0 after the first 7 weeks of the project. We can also appreciate how the recommendations generated by using FuS-KG led to a better drop of the detected undesired events compared to the ones generated by using HeLiS. A positive result has been obtained also for the DAY-Rules and the WEEK-Rules. However, we can observe how for the DAY-Rules the blue and red lines remained close for the entire timespan and, even if the FuS-KG group obtained a better drop, the improvement is not significant. This is a point of attention that triggered a more in-depth analysis of the data by combining the results observed for the MEAL-Rules with the ones observed for the DAY-Rules since it is expected that an improvement in the quality of single meals, an improvement should be observed also for the entire day. Instead, the fact that MEAL-Rules are more focused on food quantity and DAY-Rules are more focused on food categories highlighted how users found more easy-to-follow guidelines about the amount of food to consume during a single meal instead of distributing food categories appropriately during the entire day. This result will drive the analysis of how to refine the recommendations associated with DAY-Rules.

For completeness, we report in Table 2 the drop values of the observed undesired events at the beginning and at the end of the 7-week timespan of the living lab.

	FuS-KG Group	HeLiS Group	Control Group
QB-Rules	**85.06%**	76.63%	50.00%
DAY-Rules	**63.48%**	62.18%	41.98%
WEEK-Rules	**53.02%**	40.12%	6.68%

Table 2
Drop of violations at the end of the observation period. The first column lists the type of rules observed, while the second, third, and fourth columns indicate the drops observed within the FuS-KG, HeLiS, and Control groups, respectively. The most significant drop in violations occurs with the more frequent rules.

6.2. Lessons Learned

The integration of FuS-KG within the Salute+ platform and the large-scale living lab we ran, provided interesting insights, summarized in the four major lessons learned reported below, that can drive future enhancement of the overall solution.

Reasoning effectiveness. The effectiveness of the reasoner was one of the most sensitive aspects of our living lab given the unpredictability of the number of contemporary reasoning tasks that could be launched. However, the work performed on the optimization of rules design and rules evaluation schedule allowed us to maintain the time for each reasoning task below 6 seconds making the interaction with the users acceptable for a real-time context. The rule analysis and optimization process was performed in two steps. In the first one, we designed a few complex rules for covering all monitoring activities. On the one hand, we were able to cover several constraints with one rule. But, on the other hand, the computational time required for evaluating these rules was high. Hence, in the second one, we opted for splitting the rules in more simpler ones and, at the same time, scheduling their evaluation depending on their timing property. This strategy led to an improvement in the overall reasoning performance and allowed us to have easier control of the overall reasoning process (exactness of the Violation instances, debugging operations, etc.). In the scenario addressed by the current deployment of Salute+, reasoning operations are performed on sets of triple describing only users' specific events.

User perception about personalization. The second lesson is related to the actual perception that the users involved in our large-scale living lab had about the personalization capabilities of the proposed solution. To this extent, we organized a focus group at the end of the living lab with a subset of the users who participated in it. We decided on the size of 30 users for the focus group. The users were selected based on their different levels of violation drops. This way, it was possible to interview users registered a different levels of adherence to the guidelines. The aim was to collect qualitative feedback about the personalization perception of the generated recommendation by asking the users when the system succeeded and when it can be improved concerning the content of the recommendations. Besides the generally positive feedback received, during the discussion, we discovered that several users perceived the combination of some rules as very hard to follow and the Salute+ platform has been perceived as not very effective in explaining appropriately why such rules should be followed. All users who reported this issue agreed that the Salute+ system should be able to detect scenarios in which some of the rules cannot be followed by users and to automatically the user profile accordingly. This suggestion will be discussed with the domain experts to better understand, for instance, if a priority mechanism on rules can be integrated to discriminate scenarios in which the Salute+ system should discard some *UndesiredEvent* individuals.

Abandon rate. Further considerations can be made about the abandon rate of the system, a too-pushy notification system could have a high abandon rate. In our case, the percentage of the users who used the Salute+ mobile application for the entire monitoring period (i.e., seven weeks) was 87% and no complaints about the notifications were raised during the focus group. A common request to increase the engagement of the application was the possibility of better exploiting the geographical information that can be acquired through smartphone sensors. This information was considered relevant for motivating people to change habits within some real-life situations, for example not to stop at a vendor machine during a walk. Suggested examples of the exploitation of geographical information include the possibility of sending alerts about close healthy nutrition shops, restaurants cooking recipes that are compliant with users' goals, sports events related to preferred users' habits, etc. These suggestions will lead the next version of the personalization component of Salute+ to improve the perception that the system is providing enhanced real-time support to users.

Long-term adoption. At the end of the large-scale living lab, a discussion about the long-term adoption of the Salute+ was necessary with a focus on FuS-KG given the effort required to keep the KG updated and the vision of making it a reference point for the research community. We discussed this aspect through the analysis of three main perspectives: (i) the *availability and reusability* of FuS-KG; (ii) the *sustainability* plan; and, (iii) the *maintenance* plan.

Concerning availability and reusability, FuS-KG is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 [15]. It is available for download from the FuS-KG website [16] and, addi-

[15] https://creativecommons.org/licenses/by-nc-sa/4.0/
[16] https://w3id.org/fuskg

tionally, we have created a GitHub repository [17] to ease version control and issue tracking of the resource. The choice of publishing the KG open-source is to foster its adoption within the community enabling also the collection of feedback about its conceptualization to refine it. Moreover, by increasing its adoption, sustainability will benefit as well.

Concerning sustainability, as mentioned in the previous section, the presented ontology is the result of collaborative work between several experts in the context of the framework Trentino Salute 4.0. The main goals of this framework are to ‘‘combine efforts of employers, employees, and society to improve the mental and physical health and well-being of citizens’’, which is a long-term objective aligned with the Sustainable Development Goal 3 (i.e., ‘‘Ensure healthy lives and promote well-being for all at all ages’’) of the United Nations, and it aims at preventing the onset of chronic diseases related to an incorrect lifestyle through organizational interventions directed to citizens. The overall sustainability plan for the continuous update and expansion of the FuS-KG ontology is granted by this framework since it is considered a strategic asset within the AI Strategy of the Trentino Local Government.

Finally, concerning the maintenance perspective, we opted to create a collaborative environment on GitHub enabling the research community to collaborate in refining and expanding FuS-KG. This way, through the *Issues* facility all community members may open new discussions concerning specific aspects related to FuS-KG ranging from integrating new information sources to suggesting novel modeling patterns.

7. Conclusions and Future Work

In this paper, we presented the modeling pathway of FuS-KG and how it has been integrated into the Salute+ platform representing a real-world solution aiming to support the promotion of healthy lifestyles to citizens. We discussed the role of FuS-KG within the Salute+ platform and we presented the results observed within a large-scale living lab in the Trentino territory involving more the 4,200 users. Results demonstrated the possibility of adopting the system in real-world scenarios, and the reported lessons learned provide insights for future developments to improve overall efficiency, thus allowing the deployment of the Salute+ platform in more challenging environments.

References

[1] L. P. Fried, K. Bandeen-Roche, P. Chaves, B. A. Johnson, et al., Preclinical mobility disability predicts incident mobility disability in older women, Journals of Gerontology-Biological Sciences and Medical Sciences 55 (2000) M43.

[2] G. Onder, B. W. Penninx, L. Ferrucci, L. P. Fried, J. M. Guralnik, M. Pahor, Measures of physical performance and risk for progressive and catastrophic disability: results from the women's health and aging study, The Journals of Gerontology Series A: Biological Sciences and Medical Sciences 60 (2005) 74–79.

[3] F. D. Wolinsky, D. K. Miller, E. M. Andresen, T. K. Malmstrom, J. P. Miller, Further evidence for the importance of subclinical functional limitation and subclinical disability assessment in gerontology and geriatrics, The Journals of Gerontology Series B: Psychological Sciences and Social Sciences 60 (2005) S146–S151.

[4] I. LI, G. MS, Capturing and classifying functional status information in administrative databases, Health Care Financ Rev 3 (2003) 61–76. URL: https://pubmed.ncbi.nlm.nih.gov/12894635/.

[5] E. Nelson, B. Conger, R. Douglass, D. Gephart, J. Kirk, R. Page, A. Clark, K. Johnson, K. Stone, J. Wasson, et al., Functional health status levels of primary care patients, Jama 249 (1983) 3331–3338.

[6] D. R. Calkins, L. V. Rubenstein, P. D. Cleary, A. R. Davies, A. M. Jette, A. Fink, J. Kosecoff, R. T. Young, R. H. Brook, T. L. Delbanco, Failure of physicians to recognize functional disability in ambulatory patients, Annals of internal medicine 114 (1991) 451–454.

[17] https://github.com/IDA-FBK/FuS-KG

[7] D. R. Calkins, L. V. Rubenstein, P. D. Cleary, A. R. Davies, A. M. Jette, A. Fink, J. Kosecoff, R. T. Young, R. H. Brook, T. L. Delbanco, Functional disability screening of ambulatory patients, Journal of General Internal Medicine 9 (1994) 590–592.

[8] E. J. Cassell, Doctoring: The nature of primary care medicine, Oxford University Press, USA, 2002.

[9] L. P. Fried, Y. Young, G. Rubin, K. Bandeen-Roche, W. I. C. R. Group, et al., Self-reported preclinical disability identifies older women with early declines in performance and early disease, Journal of clinical epidemiology 54 (2001) 889–901.

[10] B. J. Wakefield, J. E. Holman, Functional trajectories associated with hospitalization in older adults, Western Journal of Nursing Research 29 (2007) 161–177.

[11] G. L. Alexander, M. Rantz, M. Skubic, M. A. Aud, B. Wakefield, E. Florea, A. Paul, Sensor systems for monitoring functional status in assisted living facility residents, Research in Gerontological Nursing 1 (2008) 238–244. URL: https://journals.healio.com/doi/abs/10.3928/19404921-20081001-01. doi:10.3928/19404921-20081001-01. arXiv:https://journals.healio.com/doi/pdf/10.3928/19404921-20081001-01.

[12] S. Michie, R. N. Carey, M. Johnston, A. J. Rothman, M. de Bruin, M. P. Kelly, L. E. Connell, From Theory-Inspired to Theory-Based Interventions: A Protocol for Developing and Testing a Methodology for Linking Behaviour Change Techniques to Theoretical Mechanisms of Action, Annals of Behavioral Medicine 52 (2017) 501–512. URL: https://doi.org/10.1007/s12160-016-9816-6. doi:10.1007/s12160-016-9816-6. arXiv:https://academic.oup.com/abm/article-pdf/52/6/501/27619577/s12160-016-9816-6.pdf.

[13] K. R. T. Larsen, S. Michie, E. Hekler, B. Gibson, D. Spruijt-Metz, D. Ahern, H. Cole-Lewis, R. Ellis, B. Hesse, R. Moser, J. C. Yi, Behavior change interventions: the potential of ontologies for advancing science and practice, Journal of Behavioral Medicine 40 (2016) 6–22.

[14] J. Weisz, M. Y. Ng, S. Bearman, Odd couple? reenvisioning the relation between science and practice in the dissemination-implementation era, Clinical Psychological Science 2 (2014) 58 – 74.

[15] N. Phan, D. Dou, H. Wang, D. Kil, B. Piniewski, Ontology-based deep learning for human behavior prediction with explanations in health social networks, Information Sciences 384 (2017) 298–313. URL: https://www.sciencedirect.com/science/article/pii/S0020025516306090. doi:https://doi.org/10.1016/j.ins.2016.08.038.

[16] T. W. Bickmore, D. Schulman, C. L. Sidner, A reusable framework for health counseling dialogue systems based on a behavioral medicine ontology, Journal of Biomedical Informatics 44 (2011) 183–197. URL: https://www.sciencedirect.com/science/article/pii/S1532046411000025. doi:https://doi.org/10.1016/j.jbi.2010.12.006.

[17] D. Heckmann, T. Schwartz, B. Brandherm, M. Schmitz, M. Wilamowitz-Moellendorff, Gumo - the general user model ontology, 2005, pp. 428–432.

[18] P. Kikiras, V. Tsetsos, S. Hadjiefthymiades, Ontology-based user modeling for pedestrian navigation systems, in: ECAI 2006 Workshop on Ubiquitous User Modeling (UbiqUM), Riva del Garda, 2006.

[19] A. Katifori, M. Golemati, C. Vassilakis, G. Lepouras, C. Halatsis, Creating an ontology for the user profile: Method and applications, in: RCIS, 2007.

[20] M. Dragoni, T. Bailoni, R. Maimone, C. Eccher, Helis: An ontology for supporting healthy lifestyles, in: International Semantic Web Conference (2), volume 11137 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 53–69.

[21] M. Dragoni, M. Rospocher, T. Bailoni, R. Maimone, C. Eccher, Semantic technologies for healthy lifestyle monitoring, in: The Semantic Web – ISWC 2018: 17th International Semantic Web Conference, Monterey, CA, USA, October 8–12, 2018, Proceedings, Part II, Springer-Verlag, Berlin, Heidelberg, 2018, p. 307–324. URL: https://doi.org/10.1007/978-3-030-00668-6_19. doi:10.1007/978-3-030-00668-6_19.

[22] C. Shimizu, K. Hammar, P. Hitzler, Modular ontology modeling, Semantic Web 14 (2023) 459–489. URL: https://doi.org/10.3233/SW-222886. doi:10.3233/SW-222886.

[23] M. Fernández-López, A. Gómez-Pérez, N. Juristo, Methontology: from ontological art towards ontological engineering, in: Proc. Symposium on Ontological Engineering of AAAI, 1997.

[24] H. S. Pinto, S. Staab, C. Tempich, DILIGENT: towards a fine-grained methodology for distributed,

loosely-controlled and evolving engineering of ontologies, in: R. L. de Mántaras, L. Saitta (Eds.), Proceedings of the 16th Eureopean Conference on Artificial Intelligence, ECAI'2004, including Prestigious Applicants of Intelligent Systems, PAIS 2004, Valencia, Spain, August 22-27, 2004, IOS Press, 2004, pp. 393–397.

[25] M. C. Suárez-Figueroa, NeOn methodology for building ontology networks: specification, scheduling and reuse, Ph.D. thesis, Technical University of Madrid, 2012. URL: http://d-nb.info/1029370028.

[26] M. Dragoni, I. Donadello, C. Eccher, Explainable AI meets persuasiveness: Translating reasoning results into behavioral change advice, Artif. Intell. Medicine 105 (2020) 101840. URL: https://doi.org/10.1016/j.artmed.2020.101840. doi:10.1016/J.ARTMED.2020.101840.

[27] M. Dragoni, T. Bailoni, R. Maimone, C. Eccher, Helis: An ontology for supporting healthy lifestyles, in: D. Vrandecic, K. Bontcheva, M. C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L. Kaffee, E. Simperl (Eds.), The Semantic Web - ISWC 2018 - 17th International Semantic Web Conference, Monterey, CA, USA, October 8-12, 2018, Proceedings, Part II, volume 11137 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 53–69. URL: https://doi.org/10.1007/978-3-030-00668-6_4. doi:10.1007/978-3-030-00668-6_4.

[28] G. Turconi, C. Roggi, Atlante fotografico alimentare. Uno strumento per le indagini nutrizionali, EMSI, 2007.

[29] P. Hitzler, A. Gangemi, K. Janowicz, A. Krisnadhi, V. Presutti (Eds.), Ontology Engineering with Ontology Design Patterns - Foundations and Applications, volume 25 of *Studies on the Semantic Web*, IOS Press, 2016.

[30] F. Corcoglioniti, M. Rospocher, M. Mostarda, M. Amadori, Processing billions of RDF triples on a single machine using streaming and sorting, in: ACM SAC, 2015, pp. 368–375.

[31] A. Harris, J. C. McGregor, E. Perencevich, J. Furuno, J. Zhu, D. E. Peterson, J. Finkelstein, Position paper: The use and interpretation of quasi-experimental studies in medical informatics, Journal of the American Medical Informatics Association : JAMIA 13-1 (2006) 16–23.