A Modular Strategy to Build a Large-Scale Functional Status Knowledge Graph

Mauro Dragoni[1], Tania Bailoni[1], Gabriele Sacco[1,2] and Loris Bozzato[1]

[1]Fondazione Bruno Kessler, Via Sommarive 18, Povo, Trento, Italy
[2]Free University of Bolzano, Bolzano, Italy

Abstract
Functional Status Information (FSI), which contains data on activity performance, social role involvement, and environmental and individual variables that influence well-being and quality of life, depicts physical and mental wellness at the level of the complete person. To effectively care for people with chronic diseases, multi-morbidity, and disabilities as well as to fulfill the demands of an aging global population, it is imperative to collect and analyze this data. Personal knowledge graphs (PKGs) are a useful method for displaying all data pertaining to a person's FSI in a thorough and organized manner, as well as for reasoning over them to create custom coaching solutions assisting people in everyday life for leading a healthy lifestyle. We outline the development approach of the FuS-KG, a PKG to make possible the construction of AI-enabled systems aiming to raise people's awareness of their own functional state. In particular, we emphasize the modular structure of FuS-KG designed for enhancing the scalability, management, understandability, and reuse of such a PKG.

Keywords
Knowledge Graphs, Schema Modularity, Digital Health

1. Introduction

Ontology modularization addresses the challenge of identifying a specific portion of an existing ontology for reuse. It is recognized as a fundamental principle for constructing high-quality ontologies, as stated in the guidelines of Ontology Engineering good practices [1]. Ontology developers can selectively incorporate concepts and relationships that are applicable to the specific application they are creating the ontology for.

In this paper, we present how we created a modular personal knowledge graph for representing and processing functional status information (FSI), called FuS-KG. This work started from our previous experience with both the creation of a virtual coach solution aimed at supporting people to follow healthy lifestyles [2] and the underlying ontology, called HeLiS (Healthy Lifestyle Support ontology), exploited to model the conceptual knowledge used by the system.

2nd Workshop on Modular Knowledge, 9th Joint Ontology Workshops (JOWO 2023), co-located with FOIS 2023, 19-20 July, 2023, Sherbrooke, Québec, Canada
dragoni@fbk.eu (M. Dragoni); tbailoni@fbk.eu (T. Bailoni); gsacco@fbk.eu (G. Sacco); bozzato@fbk.eu (L. Bozzato)
0000-0003-0380-6571 (M. Dragoni); 0000-0003-4360-566X (T. Bailoni); 0000-0001-5613-5068 (G. Sacco); 0000-0003-1757-9859 (L. Bozzato)

By keeping such an experience in mind, we relied on the modularization principles described in [3], and reported below, to create the modules of FuS-KG:

- *Scalability.* This principle concerns the scalability of reasoning operations. It is widely understood that, in general, the performance of reasoners degrades as the size of the ontology grows. Thus, there is a motivation to reduce the size of the ontology that needs to be reasoned over to that which is necessary. This aspect is very relevant in the scenarios in which FuS-KG is expected to be used since the reasoning task is expected to be run in real time on several batches of user-generated data. Moreover, the data used by the reasoner may relate to different user's perspectives (e.g., nutrition, activities) and reasoning operations performed on different modules will enhance the efficiency of the whole system. Then, the scalability issue also concerns the evolution of the ontology, the aim being to localize the change within an ontology module.
- *Complexity Management.* With human-designed ontologies, it becomes increasingly difficult to control the accuracy of the ontology. Ontology modularization allows the designer to focus just on the relevant portion of the ontology. In the management of the FuS-KG schema, the adoption of modularity allows for easing the modeling activities given the different clinical experts involved during the schema conceptualization.
- *Understandability.* Intuitively smaller modules are easier to understand than larger ones. This principle may be merged with the previous one even if the literature splits them. Indeed, considerations are very similar since, like in the previous point, the adoption of a modular view in FuS-KG allows the experts to better understand the content of the knowledge graph and the overall conceptualization. Indeed, the non-modular version of FuS-KG was not clear to the experts since they had to navigate through non-domain knowledge.
- *Reuse.* This is common practice in Software Engineering and Ontology Engineering would benefit from such an approach. This goal emphasizes the need for mechanisms to produce modules in such a way that increases their chances of being reused; i.e., they only contain what is relevant and useful. The reuse principle has been integrated into FuS-KG given the need of linking, when necessary, knowledge collected in one module with some concept from other modules. For example, in the *Barrier* module, nutritional barriers may be linked with one or more food categories defined within the *Food* module.

By following these principles we created the first modular version of FuS-KG, presented in Section 3. We also performed a preliminary evaluation of the whole FuS-KG (Section 4). We conclude the work with the main insights that will drive future activities (Section 5).

2. Related Work

Understanding the functional status of a person is important for developing accurate interventions and providing services for improving their health status as well as maximizing their functional independence to be able to perform daily activities well and be healthy. Managing health is a complex challenge and often the health care system's scarce resources prevent providing adequate support and care, especially to people with chronic and disability conditions [4]. In this paper, it is out of scope to present a survey on FSI methodologies given the focus on how we modularized FuS-KG in order to ease the building of AI-based applications supporting the monitoring and treatment of behavior change treatments toward the improvement of people's FSI. However, the interested reader may find further information on FSI monitoring and treatment in [5, 6, 5]

Concerning the literature on ontology modularity, we summarize below the most prominent and relevant approaches that paved the way for the work we performed about the building of FuS-KG modules. Stuckenschmidt and Klein [7] present a method for automatically partitioning ontologies based on the structure of the class hierarchy. The underlying assumption of the approach is that dependencies between concepts can be derived from the structure of the ontology. The dependencies are based on the representation language but include features such as subclass relations between concepts.

Cuenca Grau et al. [8] address the problem of partitioning an OWL ontology (O) into an σ-connection (σ). These "link relations" allow connections to be drawn between the different partitions, as such reasoning can be done on each partition individually or reasoning can be done over a combination of linked partitions.

d'Aquin et al. [9] describe an ontology module extraction technique that is integrated into the larger process of knowledge selection. Knowledge selection aims to dynamically retrieve the relevant components from online ontologies to automatically annotate a web page that is currently being viewed in a web browser.

Doran et al. [10] tackle the problem of ontology module extraction from the perspective of an Ontology Engineer wishing to reuse part of an existing ontology. The approach aims to extract an ontology module about a single user-supplied concept that is self-contained, concept centered, and consistent.

Noy and Musen [11] define the notion of traversal view extraction, which defines an ontology view. An ontology view is analogous to an ontology module because it encapsulates a subset of the original ontology. As such, this technique can be considered an ontology module extraction technique.

Seidenberg and Rector [12] developed a technique specifically for extracting an ontology module from the GALEN [1] [2] medical ontology. However, the core of the technique is generic and can be applied to other ontologies.

3. FuS-KG Modules

The personal health knowledge graph described in this paper has been built by starting from the experience of the modeling of the HeLiS [13][3] ontology, a state-of-the-art conceptual model for supporting healthy lifestyles. It defines the dietary and physical activity domains together with entities that model concepts concerning users' profiles and the monitoring of their activities. Here, we provide an overview of the modules we created for building FuS-KG starting from the core ontology (Figure 1 shows the main concepts defined within the core module of FuS-KG) by providing a summary of the seven modules integrated into the current version of the ontology.

[1] http://www.co-ode.org/galen/
[2] http://www.opengalen.org/index.html
[3] http://w3id.org/helis

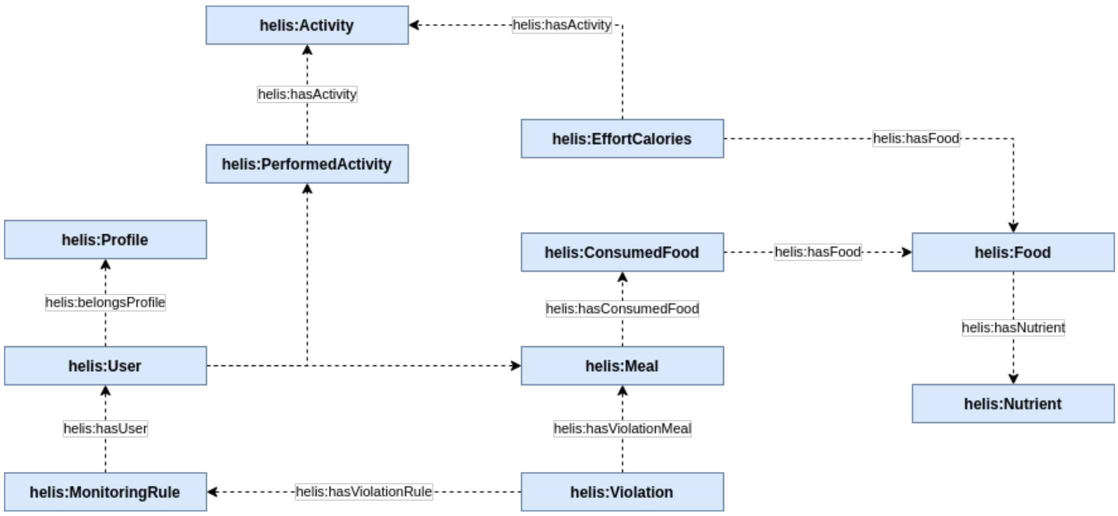

Figure 1: The main concepts of the FuS-KG core ontology.

The Food Module The *Food* module includes two macro-groups of entities descending from *BasicFood* and *Recipe* concepts. Instances of the *BasicFood* concept describe foods for which micro-information concerning nutrients (carbohydrates, lipids, proteins, minerals, and vitamins) is available, while instances of the *Recipe* concept describe the composition of complex dishes (such as *Lasagna*) by expressing them as a list of instances of the *RecipeFood* concepts. This concept reifies the relationships between each *Recipe* individual, the list of *BasicFood* it contains, and the amount of each *BasicFood*. Besides this dual classification, instances of both *BasicFood* and *Recipe* concepts are more fine-grained categorized. Concerning the number of individuals, currently, the HeLiS ontology contains 986 individuals of type *BasicFood* and 4408 individuals of type *Recipe*.

Entities subsumed by the *BasicFood* concepts are the range of the *hasMonitoredEntityType* object property linking each individual of the *MonitoringRule* concept (described below) with a specific category of food. Also, all the instances of *BasicFood* and *Recipe* concepts are within the domain of the *hasPositiveProperty* and *hasNegativeProperty* data properties having as range a string value.

Besides the food-related concepts, the classification of nutrients is also defined. The *Nutrient* concept subsumes 81 different types of nutrients properly categorized. Nutrients are instantiated through individuals describing a specific amount of a nutrient. Then each *BasicFood* is linked to all the necessary nutrients' individuals through the *hasNutrient* object property.

The Activity Module The second module relates to physical activities. The *PhysicalActivity* concept subsumes 21 subclasses representing likewise physical activity categories and a total of 856 individuals each one referring to a different kind of activity. For each activity, we defined datatype properties providing the number of calories consumed in one minute for each kilogram of weight and the MET (Metabolic Equivalent of Task) value expressing the energy cost of the activity. MET values allow to split activities in *LightActivity* (MET < 3), *ModerateActivity* (MET $[3, 6]$), and *VigorousActivity* (MET > 6).

The Monitoring Module The *Monitoring* module represents concepts concerning the monitoring of users' behaviors and it contains three main sub-concepts: *MonitoringRule*, *Violation*, and *TemporalEvent*. *MonitoringRule* instances describe the parameters defining how users should behave when adhering to health goals (aka "rules"). While *Violation* instances contain the results of reasoning activities exploited for generating users' advice and recommendations. In particular, the content of each *Violation* instance is computed according to the user data that triggered the violation. The *TemporalEvent* concept defines entities used for representing specific moments or delimited timespans to which the data to analyze refers.

Concepts subsumed by the *MonitoringEntity* one are responsible for modeling the knowledge enabling the monitoring of users' behaviors. Here, we can appreciate five concepts: *MonitoringRule*, *Violation*, *Profile*, *Goal*, and *Interval*. The *MonitoringRule* concept provides a structured representation of the parameters inserted by the domain experts for defining how users should behave. First of all, it is necessary to determine the entities affected by the monitoring rule, and the time period to be considered during the rule evaluation process. This information is provided through two object properties and one annotation property: *monitoredEntity* (e.g., *Corn* or *Walk*) and *monitoredEntityType* (e.g., *Food* or *Activity*) object properties; while the time period is provided through the *timing* annotation property that may contain the URI of the *Timespan* concept. *MonitoringRule* instances are the directives that can be exploited by a reasoner for analyzing user data. The content of the *command* datatype property specifies how the reasoner has to behave when it analyzes data of type *monitoredEntity*. The *command* is accompanied by the *hasOperator* datatype property that specifies the kind of comparison that the reasoner has to make with respect to the value/s specified through the *hasMonitoredValue* datatype property or the *hasMonitoredInterval* object property. In the first case, a classic comparison is performed between the provided data and the values contained in the monitoring rule, while, the second case indicates to the reasoner that the value specified in the rule is not a fixed value, but an interval. If the second case occurs, the reasoner will get the *ValueInterval* object linked through the *hasMonitoredInterval* object property and will check the value of the provided data with the interval specified by the *lowerBound* and *upperBound* datatype properties associated with the *ValueInterval* object.

Violation instances describe the results of the reasoning activities and they can be used by third-party applications. The content of each *Violation* instance is computed according to the user data that triggered the violation. Information materialized at runtime is contained in the *hasViolationHistory* and *hasViolationLevel* datatype properties. The former contains the number of times the *MonitoringRule* associated with the generated violation has been already triggered by the user. The latter represents the severity of the violation. Indeed, the knowledge base contains a set of pre-modeled intervals representing different levels of violations, expressed in terms of percentage with respect to the monitored value defined within the rule. When a rule is violated, the reasoner queries the violation intervals for knowing the level of the generated violation. Finally, the *hasTimestamp* datatype property contains the timestamp in which the violation instance is generated.

The *TemporalEvent* concept, as mentioned, represents specific moments or delimited timespans to which the data to analyze refers. These concepts are used in two ways. First, when users provide data concerning food consumption, these data have to be associated with a specific temporal event that, in the case of food consumption, is the *Meal* concept. In turn, the *Meal* concept subsumes other concepts defining specific kinds of meals (i.e., *Breakfast*, *Snack*, *Lunch*, and *Dinner*). Second, the other descendant of the *TemporalEvent* concept is the *Timespan* one. Instances of the concepts that subsume *Timespan* are used for driving the data selection and reasoning operations to a specific portion of data.

The User Module The *User* module contains the conceptualization of user information and it enables the representation of all users' events (i.e., consumed foods and performed physical activities) and the association of each violation to the corresponding user. Users' events are represented via the *Meal*, *ConsumedFood*, and the *PerformedActivity* concepts. The last two concepts are reified relations enriched with attributes for representing the facts that a user consumed a specific quantity of food or performed an activity for a specific amount of time.

This concept subsumes the conceptualization of information that a user can provide, such as food consumption and performed activities, and also links them with the possible violation that can be generated after their analysis. Concerning the representation of users' activities and personalized information, we modeled the *ConsumedFood* and the *PerformedActivity* concepts. Both concepts are used as the reification of the fact that a user consumed a specific quantity of some food or performed an activity for a specific amount of time. In the first case, every *Meal* is associated with a list of *ConsumedFood* through the *hasConsumedFood* object property. While, in the second case, instances of the *PerformedActivity* concept associate a user with the amount of time they spent in performing a specific activity. Here, we did not use a concept for grouping a list of activities a user is routinely doing (including sleeping), thus a further concept for grouping activities is useless. Then, we included also the possibility of providing user-specific information representing the energy consumption equivalence, e.g., how much a specific user has to run for burning 100 grams of pasta. This information is represented by instantiating the *EffortCaloriesEquivalent* concept (EffortCalories in Figure 1). Then, instances of the *User* concept are used as objects for the *hasViolationUser* object property defined on the *Violation* concept (described below).

The Enablers Module The *Enablers* module contains the main concepts enabling a user to start a behavior change process. This module has been built by starting from two of the main references available in the field indicated by the domain experts [14, 15]. We defined four main concepts, namely *Intervention*, *Treatment*, *Strategy*, and *Technique*.

The *Intervention* concept follows the definition provided in [15] and it refers to a single action performed during a *Treatment*. As an example, let us consider a behavior change scenario where a patient affected by diabetes has to monitor their glycemic index after each meal and provide the observed value into a mobile application. The reminder to do this action is an instance of the *Intervention* concept.

The *Treatment* concept is defined as the unfolding of every *Intervention* performed to allow users to achieve their aim. For example, the set of *Intervention*s performed by an AI-enabled system to persuade a patient to follow the Mediterranean Diet is a *Treatment* [16].

The *Strategy* concept subsumes the five main strategies that can be implemented during a behavior change process: *Education*, *GoalAndPlanning*, *Feedback*, *Monitoring*, and *MotivationalEnhancement*. The *Education* concept models the aim to increase the user's understanding of their past and current state and of the steps required to achieve the future state (e.g., to provide information and/or instruction for behaving in a proper way). The *GoalAndPlanning* concept refers to future planning to achieve desired future states (e.g., activity scheduling and/or setting tasks of progressively greater difficulty). The *Monitoring* concept defines the action of recording past or current user's states (e.g., current nutritional behaviors and/or activity events). The *Feedback* concept models the information on current and past states provided to the user about their condition and/or actions. The meaning of the *Feedback* concept may also overlap with other behavior intervention components, such as *MotivationEnhancement* in a scenario where, for example, feedback may provide information about progress and may also increase or decrease motivation. Finally, the *MotivationEnhancement* concept refers to interventions that increase the likelihood that the user will engage in specific behaviors related to treatment goals or usage of the application in the future. Each instance of the *Strategy* concept has to be associated with the instance of the *Treatment* concept adopting it. Such an association can be done by instantiating the *adoptsStrategy* object property.

The fourth main concept is *Technique*, meaning an observable, replicable, and irreducible component of a *Strategy* used within a *Treatment* designed to alter or redirect causal processes that regulate behavior. A *Technique* is an "active ingredient" (e.g., feedback, self-monitoring, reinforcement) of a *Strategy* and it can be used alone or in combination, and in a variety of formats. Within the proposed FuS-KG, we defined 19 types of techniques subsuming the *Technique* concept. We refer the reader to check the *Enablers* module for the details. Moreover, each instance of the *Technique* concept has to be associated with the instance of the *Strategy* concept adopting it. Such an association can be done by instantiating the *isUsedBy* object property.

The Barriers Module The Barriers ontology module is composed of two main branches: (i) the classification of the barriers, and (ii) the representation of the different states of change.

The *Barrier* concept is the root concept of the first branch and it subsumes six macro-categories of barriers. The *EnvironmentBarrier* refers to the hindrances related to performing an action due to obstacles connected to the circumstances in which the action itself takes place. For example, they could relate to the weather (like unfavorable climatic conditions), to money (like the cost of the equipment needed), and to security issues (like the lack of safety). *HealthBarrier* concerns the presence of some disease preventing the performance or completion of a specific action. This concept enables the possibility of importing external medical knowledge bases (e.g., the UMLS). This way, barriers are connected with medical knowledge that can be exploited at reasoning time (such as asthma, chest pain, etc.). The *PersonalBarrier* concept represents barriers associated with real-life situations (e.g., job conditions) that obstruct the performance of specific actions. Then, the *PhysicalBarrier* and *PsychologicalBarrier* concepts are related to hindrances given by physical pains (e.g., knee injury) or emotional status (e.g., fear) that block a person from performing specific actions. Finally, the *SocialBarrier* concept mainly refers to a possible lack of support from people close to the patients (e.g., parents, friends, etc.). The *Barrier* concept is associated with the *Sign* concept specifying a condition that is true in relation to that barrier. In fact, the *Sign* concept defines specific circumstances (both environmental and personal) like the fact that it's raining, that a person has a headache, or that a person has a full-time job.

The second branch consists of the abstract representation of the Transtheoretical Model of change (TTM) [17]. TTM describes the different stages of change that an individual can be in and is used by clinicians for supporting the behavior change process. The main concepts we defined are *StateOfChange*, which is the root concept of this branch, and then the six stages in which a *Patient* can be: *PreContemplation*, *Contemplation*, *Preparation*, *Action*, *Maintenance*, and *Termination*. Moreover, we defined the property *hasBehavior* that is used as a reification of the status in which a *Patient* is during a specific *Timespan*.

The Arguments Module The arguments ontology module aims at supporting efficient and effective dialogues for motivating behavior change. It has been developed following the model described in [18], and it is structured in a way that helps the formulation of persuasive dialogues [19] aimed at motivating a user/patient toward a healthier lifestyle (or a particular lifestyle goal). The module collects various arguments, which are different types of sentences, that are used to model dialogues providing beliefs a person may have concerning healthcare issues and the appropriate responses motivating behavior change. Therefore, each argument can have various properties that help to define its type, function, topic, context, utility, and the healthcare problem or solution it refers to. Furthermore, the concepts defined in the arguments module are *Argument*, *Concern_category*, *Context_type*, *Functional_type*, *Ontological_type*, *Topic_type* and *Utility*.

4. Empirical Evaluation

The evaluation of the quality and correctness of FuS-KG has been conducted from two perspectives. First, we performed an expert-based evaluation where the team that did not participate in the modeling process adopted the metrics described in [20, 21, 22, 23, 24] to verify the quality of FuS-KG: *Accuracy*, *Adaptability*, *Clarity*, *Completeness*, *Computational Efficiency*, *Conciseness*, *Consistency/Coherence*, and *Organizational fitness*. Second, we ran the OntOlogy Pitfall Scanner! (OOPS!) [25] to identify inconsistencies, pitfalls, and errors and to check whether FuS-KG met all the needs for which it has been built.

The overall *Accuracy* of FuS-KG has been judged as satisfactory. The knowledge of the domain experts was in-line with the complexity of the use of axioms. Indeed, within the FuS-KG there are not very complex axioms. Then, by considering the representation of the real world, the evaluators agreed on the correctness of the FuS-KG in describing the domain.

Concerning the *Adaptability* of the FuS-KG, the evaluators focused on the possible extension aspects. They verified that the FuS-KG can be extended and specialized monotonically. Here, the question has to be addressed from two perspectives. Firstly, concerning the extension of the FuS-KG from the content perspectives (i.e., adding new enablers, barriers, FSI-related concepts, etc.), the result was positive because any extension of the FuS-KG did not require the removal of any axiom. Secondly, concerning the representation of users' profiles, the update of the FuS-KG was not monotonic because if a user is associated with a new profile, the old association is removed. Anyway, the FuS-KG does not react negatively to these changes because its consistency is preserved.

About the *Clarity* of the FuS-KG, the evaluators agreed with the strategy decided by the modeling team about using concept labels to communicate the intended meaning of each concept and the use of definitions and descriptions of the main concepts of the FuS-KG, especially for the root concepts of each branch. Moreover, each definition has been well documented within the FuS-KG in order to make the meaning of each concept understandable by those who use the knowledge graph.

The experts agreed about the *Completeness* of the FuS-KG. However, they distinguished between the TBox and the ABox. Indeed, concerning the TBox, the evaluators agreed about the completeness of the FuS-KG and the lexical representations of the concepts. In particular, they verified that all the represented nutrients appropriately cover the health domain and that all the information needed for the realization of tools supporting a healthy lifestyle were modeled within the FuS-KG. Regarding ABox, the evaluators highlighted the necessity of including individuals concerning commercial products. This observation is interesting, especially, if we consider the possibility of developing end-user applications. Indeed, the presence of commercial products will improve overall user engagement.

In order to verify the *Computational efficiency* of the FuS-KG, we observed how it behaved within the scenario introduced in Section 1 and detailed in [2]. Indeed, the FuS-KG itself does not contain axioms representing a criticism for reasoners. On the contrary, the final aim of the FuS-KG is to be used for analyzing data provided by users. In [2], we show an example of how the FuS-KG is used and we provide statistics regarding the amount of time needed for completing the reasoning activity with respect to the dimension of elaborated data.

The evaluators judged the FuS-KG "*Concise*" because all the axioms included are relevant with regard to the targeted domain and there are no redundancies. Also, the FuS-KG has been judged "*Consistent*" and "*Coherent*". It has been judged consistent because no contradictions were found by the evaluators and coherent because the evaluators observed little bias between the documentation containing the informal description of the concepts and their formalization.

Then, concerning the *Organizational fitness*, the FuS-KG has been deployed within the organization as a web service in order to make it easily accessible by the community and potential stakeholders. A focus group has been organized with both the modeling team and the evaluation team for discussions about the adopted methodology, which was judged appropriate by considering the necessity of working in situ altogether and synchronizing the commitments of all the people involved.

Finally, the entire FuS-KG has been analyzed by the OOPS! tool to find potential pitfalls and trigger mitigation actions. Since some ontologies have been reused in the core module of FuS-KG (i.e., the HeLiS ontology), some pitfalls appeared, but all of them pointed to reused ontology entities. The pitfall record includes P04 (Creating unconnected ontology elements), P08 (Missing annotations), P11 (Missing domain or range in properties), P13 (Inverse relationships not explicitly declared), and P22 (Using different naming conventions in the ontology). However, all pitfalls that appeared in the modules for newly implemented entities during the implementation have been solved. Inconsistencies were checked with the reasoners Pellet [26] and HermiT [27] and no errors were found when running them. To ensure that each module met the quality needs, the scanning has also been performed separately on each module.

Evaluations on the whole FuS-KG and for all modules were run successfully, with good results. This ensures that FuS-KG is consistent, meets the requirements, and to the best of our knowledge, has no errors.

5. Conclusions

In this work, we presented how FuS-KG modules have been created and the rationale behind the choices we made. These modules aim to enhance the AI capabilities of coaching systems designed for supporting the monitoring of users' functional status from both effective and efficient perspectives.

As mentioned in Section 1, this work represents a first step toward the long-term achievement of having a full-fledged AI coaching system. Future efforts will be focused on three main directions. First, we plan to expand the knowledge base since ontologies are inevitably subject to constant changes. This will involve domain experts and the exploration of techniques that leverage some form of data mining able to detect hidden information from large textual data. Secondly, we plan to integrate ontology with natural language understanding (NLU) and natural language generation (NLG) components. This way, we will be able to investigate strategies about how to automatically transform natural language texts into their equivalent semantic argument-based representation, as well as, exploit the output of the reasoning process for generating effective contextual feedback. Thirdly, we plan to evaluate the system in a real-world coaching scenario. In this work, we did not provide a living lab evaluation since the FuS-KG itself cannot be evaluated without addressing the points above (i.e., integration with both NLU and NLG). We will focus on doing such integration in order to deploy the end-to-end system into real-world scenarios and to observe the effectiveness of the ontology modules presented.

Acknowledgments

We acknowledge the support of the PNRR project FAIR - Future AI Research (PE00000013) and of the PNRR project INEST - Interconnected Nord-Est Innovation Ecosystem (ECS00000043), under the NRRP MUR program funded by the NextGenerationEU

References

[1] A. Bernaras, I. Laresgoiti, J. M. Corera, Building and reusing ontologies for electrical network applications, in: W. Wahlster (Ed.), 12th European Conference on Artificial Intelligence, Budapest, Hungary, August 11-16, 1996, Proceedings, John Wiley and Sons, Chichester, 1996, pp. 298–302.

[2] M. Dragoni, M. Rospocher, T. Bailoni, R. Maimone, C. Eccher, Semantic technologies for healthy lifestyle monitoring, 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. 307–324. URL: https://doi.org/10.1007/978-3-030-00668-6_19. doi:10.1007/978-3-030-00668-6_19.

[3] H. Stuckenschmidt, C. Parent, S. Spaccapietra, Modular ontologies: concepts, theories and techniques for knowledge modularization, volume 5445, Springer, 2009.

[4] A. Baker, Crossing the quality chasm: a new health system for the 21st century, volume 323, British Medical Journal Publishing Group, 2001.

[5] 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.

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

[7] H. Stuckenschmidt, M. C. A. Klein, Structure-based partitioning of large concept hierarchies, in: S. A. McIlraith, D. Plexousakis, F. van Harmelen (Eds.), The Semantic Web - ISWC 2004: Third International Semantic Web Conference,Hiroshima, Japan, November 7-11, 2004. Proceedings, volume 3298 of *Lecture Notes in Computer Science*, Springer, 2004, pp. 289–303. URL: https://doi.org/10.1007/978-3-540-30475-3_21. doi:10.1007/978-3-540-30475-3_21.

[8] B. C. Grau, B. Parsia, E. Sirin, A. Kalyanpur, Automatic partitioning of OWL ontologies using *E*-connections, in: I. Horrocks, U. Sattler, F. Wolter (Eds.), Proceedings of the 2005 International Workshop on Description Logics (DL2005), Edinburgh, Scotland, UK, July 26-28, 2005, volume 147 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2005. URL: https://ceur-ws.org/Vol-147/21-Grau.pdf.

[9] M. d'Aquin, M. Sabou, E. Motta, Modularization: a key for the dynamic selection of relevant knowledge components, in: P. Haase, V. G. Honavar, O. Kutz, Y. Sure, A. Tamilin (Eds.), Proceedings of the 1st International Workshop on Modular Ontologies, WoMO'06, co-located with the International Semantic Web Conference, ISWC'06 November 5, 2006, Athens, Georgia, USA, volume 232 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2006. URL: https://ceur-ws.org/Vol-232/paper2.pdf.

[10] P. Doran, V. A. M. Tamma, L. Iannone, Ontology module extraction for ontology reuse: an ontology engineering perspective, in: M. J. Silva, A. H. F. Laender, R. A. Baeza-Yates, D. L. McGuinness, B. Olstad, Ø. H. Olsen, A. O. Falcão (Eds.), Proceedings of the Sixteenth ACM Conference on Information and Knowledge Management, CIKM 2007, Lisbon, Portugal, November 6-10, 2007, ACM, 2007, pp. 61–70. URL: https://doi.org/10.1145/1321440.1321451. doi:10.1145/1321440.1321451.

[11] N. F. Noy, M. A. Musen, Specifying ontology views by traversal, in: S. A. McIlraith, D. Plexousakis, F. van Harmelen (Eds.), The Semantic Web - ISWC 2004: Third International Semantic Web Conference,Hiroshima, Japan, November 7-11, 2004. Proceedings, volume 3298 of *Lecture Notes in Computer Science*, Springer, 2004, pp. 713–725. URL: https://doi.org/10.1007/978-3-540-30475-3_49. doi:10.1007/978-3-540-30475-3_49.

[12] J. Seidenberg, A. L. Rector, Web ontology segmentation: analysis, classification and use, in: L. Carr, D. D. Roure, A. Iyengar, C. A. Goble, M. Dahlin (Eds.), Proceedings of the 15th international conference on World Wide Web, WWW 2006, Edinburgh, Scotland,

UK, May 23-26, 2006, ACM, 2006, pp. 13–22. URL: https://doi.org/10.1145/1135777.1135785. doi:10.1145/1135777.1135785.

[13] 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.

[14] S. Michie, R. West, R. Campbell, J. Brown, H. Gainforth, ABC of Behaviour Change Theories, 2014.

[15] D. C. Mohr, S. M. Schueller, E. Montague, M. N. Burns, P. Rashidi, The behavioral intervention technology model: an integrated conceptual and technological framework for ehealth and mhealth interventions, Journal of medical Internet research 16 (2014) e3077.

[16] I. Donadello, M. Dragoni, C. Eccher, Explaining reasoning algorithms with persuasiveness: a case study for a behavioural change system, in: SAC, ACM, 2020, pp. 646–653.

[17] J. Prochaska, C. Redding, K. Evers, The transtheoretical model and stages of change, 2015, p. 97.

[18] L. Chalaguine, E. Hadoux, F. Hamilton, A. Hayward, A. Hunter, S. Polberg, H. W. W. Potts, Domain modelling in computational persuasion for behaviour change in healthcare, 2018. arXiv:1802.10054.

[19] I. Donadello, A. Hunter, S. Teso, M. Dragoni, Machine learning for utility prediction in argument-based computational persuasion, in: AAAI, AAAI Press, 2022, pp. 5592–5599.

[20] L. Obrst, W. Ceusters, I. Mani, S. Ray, B. Smith, The Evaluation of Ontologies, Springer US, Boston, MA, 2007, pp. 139–158.

[21] A. Gangemi, C. Catenacci, M. Ciaramita, J. Lehmann, Modelling ontology evaluation and validation, in: Y. Sure, J. Domingue (Eds.), The Semantic Web: Research and Applications, 3rd European Semantic Web Conference, ESWC 2006, Budva, Montenegro, June 11-14, 2006, Proceedings, volume 4011 of *Lecture Notes in Computer Science*, Springer, 2006, pp. 140–154.

[22] A. Gómez-Pérez, Ontology evaluation, in: S. Staab, R. Studer (Eds.), Handbook on Ontologies, International Handbooks on Information Systems, Springer, 2004, pp. 251–274.

[23] T. R. Gruber, Toward principles for the design of ontologies used for knowledge sharing?, Int. J. Hum.-Comput. Stud. 43 (1995) 907–928. URL: https://doi.org/10.1006/ijhc.1995.1081. doi:10.1006/ijhc.1995.1081.

[24] M. Grüninger, M. Fox, Methodology for the Design and Evaluation of Ontologies, in: IJCAI'95, Workshop on Basic Ontological Issues in Knowledge Sharing, April 13, 1995, 1995. URL: http://citeseer.ist.psu.edu/grninger95methodology.html.

[25] M. Poveda-Villalón, A. Gómez-Pérez, M. C. Suárez-Figueroa, Oops! (ontology pitfall scanner!): An on-line tool for ontology evaluation, Int. J. Semantic Web Inf. Syst. 10 (2014) 7–34. URL: https://doi.org/10.4018/ijswis.2014040102. doi:10.4018/ijswis.2014040102.

[26] E. Sirin, B. Parsia, B. C. Grau, A. Kalyanpur, Y. Katz, Pellet: A practical OWL-DL reasoner, J. Web Semant. 5 (2007) 51–53. URL: https://doi.org/10.1016/j.websem.2007.03.004. doi:10.1016/j.websem.2007.03.004.

[27] B. Glimm, I. Horrocks, B. Motik, G. Stoilos, Z. Wang, Hermit: An OWL 2 reasoner, J. Autom. Reason. 53 (2014) 245–269. URL: https://doi.org/10.1007/s10817-014-9305-1.