

An Ontology-based Decision Support Framework for Personalized Quality of Life Recommendations

Marina Riga^{1,2} (✉), Efstratios Kontopoulos¹, Kostas Karatzas²,
Stefanos Vrochidis¹ and Ioannis Kompatsiaris¹

¹ CERTH-ITI, Information Technologies Institute, Thessaloniki, Greece
{*mriga, skontopo, stefanos, ikom*}@iti.gr

² ISAG-EI Group, Dept. of Mechan. Engineering, Aristotle Univ. of Thessaloniki, Greece
{*mriga, kostas*}@isag.meng.auth.gr

Abstract. As urban atmospheric conditions are tightly connected to citizens' quality of life, the concept of efficient environmental decision support systems becomes highly relevant. However, the scale and heterogeneity of the involved data, together with the need for associating environmental information with physical reality, increase the complexity of the problem. In this work, we capitalize on the semantic expressiveness of ontologies to build a framework that uniformly covers all phases of the decision making process: from structuring and integration of data, to inference of new knowledge. We define a simplified ontology schema for representing the status of the environment and its impact on citizens' health and actions. We also implement a novel ontology- and rule-based reasoning mechanism for generating personalized recommendations, capable of treating differently individuals with diverse levels of vulnerability under poor air quality conditions. The overall framework is easily adaptable to new sources and needs.

Keywords: personalized decision support, ontology, OWL 2, SPIN, air quality recommendations, user profiling.

1 Introduction

In the environmental domain, it was early recognized that there is a positive correlation between the quality of the atmospheric environment and people's quality of life (QoL), acknowledging the fact that sensitive parts of the population suffer as atmospheric quality parameters become worse. Although air quality (AQ) experts could benefit from rapid advances in information technology and telecommunications, which enabled the efficient monitoring, analysis, transmission and sharing of scientific AQ data, it was only after the European Union's Directive 97/101/EC that environmental information should mandatorily become available and easily accessible to the public via official and direct communication channels.

As a first result, several urban air quality management and information systems emerged. However, the delivered content was usually limited to spatiotemporal air

quality observations, emission data, or AQ forecasts [1], with no sufficient explication of their meaning or impact to individuals' health. The provision of the aforementioned content in a way that can support end-users in decision-making activities is an integral part of the so called *environmental decision support system (EDSS)*.

The way communication of AQ information is handled in EDSSs is important for their wider acceptance. Humans perceive environmental quality on the basis of their personal interests; they need to be informed if caused events have an impact on the daily living [2-3]. Hence, systems that link the *location*, the *incident*, and the *individual* are of high interest for the end-users [4]. In this direction, relevant initiatives, such as PESCaDO¹, AirForU², Clean Air Nation³ and Air Visual⁴, demonstrate the added value of real-time, location-based AQ information and recommendation services. However, these applications do not handle user-profiles' differentiation, but produce general advice under poor AQ conditions that universally apply to sensitive people and not to the specific user who queries for decision support (DS).

Our motivation is to integrate the involved data and processes of an EDSS in a uniform, modular and user-profile centric framework. The effectiveness of this demanding task lies in the following subtasks: (i) to efficiently handle the heterogeneous and multifaceted nature of data, (ii) to adequately associate the provided experts' knowledge and rules for targeted, user-profile driven recommendation provision, and (iii) to facilitate scalability and reusability of the framework to third-party modules. In this context, we take advantage of the semantic expressiveness of ontologies to deal with the above issues. *Ontologies* are state-of-the-art Semantic Web technologies for structuring and semantically integrating heterogeneous content. Among other applications, they have been successfully adopted for covering individual parts of the decision making process [5]: (i) collecting, storing, and processing data, (ii) formulating the decision-making problem, (iii) reasoning over the data to reach decisions.

In contrast to existing ontology-based decision support systems (DSSs) that merely exploit the semantic web technologies in parts of the decision making process (see next section), our proposed approach demonstrates the extensive use of ontologies and semantic reasoning technologies by handling both the *static* (representation) and *dynamic* processes (realization, inference) of a DSS operation. The proposed schema comprises a set of ontological concepts and relations for semantically representing a primitive section of experts' knowledge and AQ dynamics, focusing on definitions of air pollutants' measurements, health risks, sensitive groups, and relevant user-profile driven recommendations. We implement a novel rule-based ontological reasoning mechanism for routing the problem of dynamic classification and new knowledge extraction to support personalized recommendation provision. To the best of our knowledge, no other DSS covers this multifaceted task as a whole through the adoption of ontologies. The proposed work comprises the operational EDSS of the hackAIR EU project [6].

¹ <http://pescado-project.upf.edu/>

² <http://newsroom.ucla.edu/releases/new-app-lets-you-check-air-quality-as-easily-as-checking-the-weather>

³ <http://www.greenpeace.org/india/Clean-Air-Nation/>

⁴ <https://airvisual.com/app>

2 Related Work

An EDSS literature review reveals various implementations, with respect to: (i) the application domain for DS (e.g. urban air quality management, extreme climate risks administration, etc.) and (ii) the technological approaches (e.g. computational intelligence methods, mathematical models, etc.). A detailed review of EDSS technologies, tools, and use cases is presented in [7-8]. In most related work items, DS is addressed for administrative purposes, involving experts or regulatory authorities as the targeted end-users. Instead, our proposed framework aims to support QoL services for the general public, with an additional strategic difference against existing implementations, the process of “translating” AQ observations into user-profile driven recommendations, for personalized guidance in severe atmospheric conditions.

Considering ontologies in the environmental domain, numerous implementations exist covering the representation of abstract, general concepts (e.g. SWEET, a modular schema with 6,000 concepts in 200 ontologies describing concepts of physical and ecological phenomena, meteorological conditions, processes, activities [9]) as well as domain-specific, applied concepts (e.g. the PESCaDO ontology for personalized environmental DS [5], the EnvO ontology for the concise description of environmental features, materials and habitats [10], the AIR_POLLUTION_Onto ontology for air pollution analysis and control [11], etc.). Inspired by their expressiveness and adaptability, we build our relevant representations as described in the next section.

Ontologies have been proposed in EDSS for different tasks: in [12] for semantic search and easy access of structured environmental data; in [13] for integrating existing local databases of environmental data as part of the Linked Open Data cloud, enabling the linking of data in an established context and the dissemination of environmental information to the masses; in [14] for facilitating the process of selecting domestic solar hot water systems according to specific criteria, and in [15] for integrating heterogeneous content from multiple environmental sources. Despite the increasing deployment of ontology-based solutions in DSS, their potential is merely exploited, either for creating a structured representation of the domain of interest, or for supporting parts of the decision making process. With our proposed framework, we demonstrate the efficient use of ontologies and their supported technologies in all the basic components of the DSS.

Similarly to our approach, within the context of the PESCaDO EU project, ontologies were used as the backbone of the proposed EDSS, supporting all phases of the decision making process [16]; nevertheless, its rules along with the reasoning module are hardcoded in the source code, resulting in a highly inflexible approach. On the other hand, our implementation pushes the usage of ontologies one step further: both the domain knowledge and the experts’ rules are developed at the ontology level, with the use of the OWL language and SPIN rules (see next section). The proposed DS framework operates as a stand-alone and uniformly developed module that can be easily adoptable by external sources, independently of their implemented technologies. It is also flexible and extensible, in terms of easy initialization of the different concepts, rules and recommendations even by non-ontology experts, by simply following the definitions of the proposed schema.

3 Proposed Approach

In computer science, an ontology is defined as a formal explicit specification of the terms and relations that describe a domain of discourse in a structured and semantically rich way [17]. The adoption of ontologies enables the understanding, sharing and reuse of information among different systems. Their capabilities fit perfectly to the task of describing and integrating heterogeneous content, and of dynamically inferring new knowledge, in a multidisciplinary field of study such as air quality. In this paper, we do not exhaustively represent the complete dynamics and facts or relevant association rules existing in the AQ domain; instead we demonstrate a practical paradigm of ontology use for real scenarios, conforming to related literature suggestions stating that citizens as end-users seek *personalized information services*, with *timely* or *in advance* AQ information provision, with respect to *their location* of interest [4].

The proposed approach handles all three basic components of a general DSS [18] through the use of ontologies: (i) the *data* component, i.e. the ontology-based schema developed for representing an excerpt of AQ domain experts' knowledge and of AQ information provision; (ii) the *model* component, which is an ontology-based and augmented with rules dynamic representation of principles that generate recommendations by combining different disciplines in the field of interest; and (iii) the *user-system interaction* component, which involves the communication of the user with the system. All of the above are more thoroughly described in the following subsections.

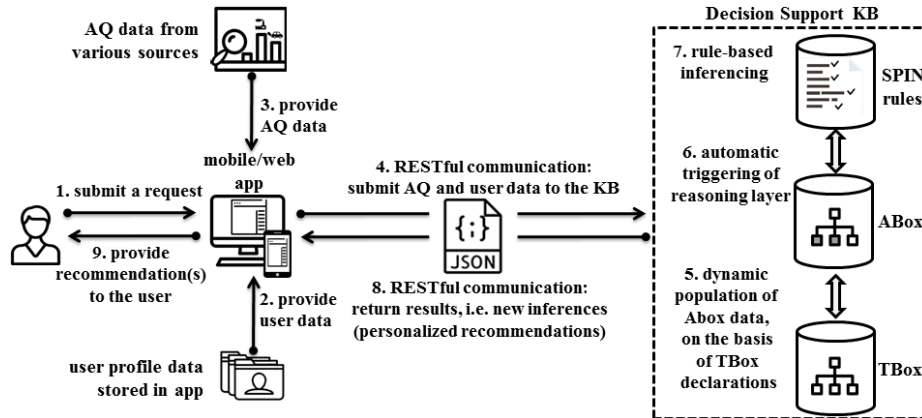


Fig. 1. Overall architecture and workflow of the proposed approach

In this context, we implemented a *3-layered* decision support knowledge base (KB) on the basis of three interconnected ontologies: (1) the *TBox*, i.e. the bottom layer that formalizes information relevant to the concepts of discourse (user profile, AQ measurements, requests and personalized recommendations) and their interrelations, (2) the *ABox*, i.e. the middle layer that is based on the schema from the TBox declarations and formalizes information relevant to membership/attribute assertions (actual users, observations, etc.), and (3) the *rules* layer, i.e. the upper layer that is based on both the aforementioned layers and formalizes the set of rules for reasoning (inference) differ-

ent levels of interpretations. A general overview of the proposed ontology-based architecture and information workflow is presented in Fig. 1

3.1 Domain Knowledge Modelling

The representational primitives utilized in ontologies are: *classes* referring to concepts or (abstract) entities that are assumed to exist in some domain of interest, *individuals*, which are instantiations (i.e. objects) of the classes, and *properties*, i.e. relationships that hold among objects. Their formal definition within the ontology typically carries information about concepts' meaning (semantics) and constraints that exist within the actual context of the domain of discourse.

To elicit the requirements that our ontology should satisfy, we followed the guidelines proposed by the *NeOn methodology* [19]. First, we defined the multifaceted *purpose of use* of our ontology, which includes the following goals: (i) to serve as an operational framework for the representation and orchestration of heterogeneous environmental, health, user profile-related data; (ii) to integrate the rules that govern air pollution and their impact to QoL, according to provided environmental and health experts' knowledge; and (iii) to support user-oriented recommendation services, with respect to personal health/user preferences (i.e. activities, daily routine, asthma, etc.) and to current AQ conditions for the location of interest.

Then, we identified the ontology's *intended users* as: (i) individuals, people with health sensitivities, those working or exercising outdoors, all those interested in receiving information about existing AQ conditions, so as to limit their exposure to hazardous conditions, or to increase their awareness about the impact of air pollution under special circumstances; (ii) public administrators or environmental experts, interested in receiving AQ information for professional reasons; (iii) technology experts, developers or ontology engineers, interested in adopting and expanding the ontology model for relevant DSSs and services.

The aforementioned definitions implicitly define the content and structure of the ontology. The latter is developed in OWL 2, a W3C standard ontology language [20]. Here, we present an excerpt of conceptualizations that are significant for adequately structuring the DS process following a user request. An indicative graph of the main classes and their relationships, developed in *Grafoo* [21], is presented in Fig. 2.

Class `Person` encapsulates those individuals' characteristics that are required for the reasoning process, specified on the basis of the available health-related advice and knowledge. We define distinct subclasses of class `Person`, with respect to the following parameters: (i) year of birth (`ChildPerson`, `ElderlyPerson` etc.), (ii) health sensitivity (`PregnantFemalePerson`, `SensitiveHealthPerson`), (iii) daily preferred activities (`SportsWalkingPerson`, `OutdoorJobPerson` etc.). Specific *ontology rules* (see next section) and *class expression axioms* (e.g. `EquivalentClasses (OutdoorJobPerson ObjectSomeValuesFrom (worksOutdoors "true"^^xsd:boolean))`) declare the underlying semantics and restrictions.

In our proposed schema, a user may be classified as an instance of more than one types (e.g. `SensitiveHealthPerson` and `SportsWalkingPerson`), unless such combinations are forbidden (e.g. `DisjointClasses (ElderlyPerson ChildPerson)`). Moreover, a `Person` instance can be associated with other linked user

profiles, enabling a *one-to-many* DS provision. An example case could be a mother with respiratory problems requesting personalized AQ recommendations for herself and her child, where separate recommendations will be produced from a single request, with respect to each profile's characteristics.

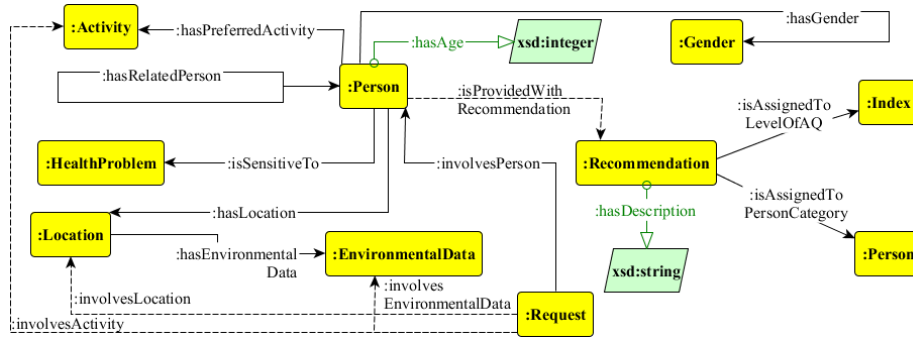


Fig. 2. An excerpt of classes and relations of the proposed ontology; properties with dashed line are inferred via rules.

Class *Activity* represents instances of indoor/outdoor activities. On the basis of rules, recommendations to the users differ according to the nature or intensity of the activity; e.g. in case of harmful AQ conditions, activities like running or biking may lead to increased oxygen uptake and should be avoided or replaced with alternatives.

Class *HealthProblem* represents health problems (e.g. respiratory, cardiovascular, circulatory diseases) that the DSS takes into account when generating recommendations with respect to users' health sensitivities, under poor air quality conditions.

Class *Location* represents the location of interest (city, country or as coordinates). The more discretization introduced in the ontology, the more complex rule-definitions are needed to support the decision making process, i.e. to propose alternative areas with less pollution or close to the area of interest.

Class *EnvironmentalData* represents environmental measurements, i.e. observations from different sources regarding air pollutants, pollen, weather or any other measurable environmental aspect that is involved in the recommendation process.

Class *Request* structures the content of a user-request for DS. An instance of that type connects all these notions (user profile characteristics, the location of interest, the existing AQ measurements, and preferred activities) that are fed into the rule-based DSS for personalized recommendation provision.

Class *Recommendation* represents messages with fixed content, as defined by environmental experts for recommendation, together with details on: (a) which type of user such messages concern, and (b) under which AQ conditions they should be informed. An instance of that type may contain the actual message to be inferred to the user. Rules defined in the respective layer of the framework handle the matching between categories of users and defined recommendations.

3.2 Reasoning over Domain Knowledge

For the semantic interpretation and inference of new knowledge, we implemented a rule-based reasoning layer that fully complies with the *SPARQL Inferencing Notation (SPIN)* framework, which is a well-established standard for representing SPARQL rules and constraints on Semantic Web models [22]. SPIN rules are linked to ontology classes and stored alongside the domain model as RDF triples, thus supporting a holistic, dynamic, semantically enriched approach which fits perfectly to the requirements of our defined DS problem.

In the proposed framework, the rule layer is implemented separately from the abstract schema (TBox) and the assertions (ABox), and the rules are grouped in two different levels. Thus, the reasoning process operates incrementally, where inferences from one step serve as input to rules of the following steps, moving from more general to more specific derivations. Such a distinction and hierarchy of rules facilitates the extensibility of the reasoning module, eliminating the intervention to high levels of rules when lower levels change, and vice versa.

In order to fully exploit SPIN's capabilities, we capitalize on the SPIN vocabulary (<http://spinrdf.org/spl>); among the available concepts, we adopt/extend three specific types: `spin:Functions`, `spin:Rule`, and `spin:MagicProperties`, which serve different functionalities (select/construct triples, reuse SPARQL blocks, etc.).

```
CONSTRUCT {
  ?person :isProvidedWithRecommendation ?personalized_rec .      #1
  ?personalized_rec a :LimitExposureRecommendation .             #2
  ?personalized_rec :hasDescription ?rec_description .           #3
  ?personalized_rec :hasDescriptionIdentifier "walking" .        #4
} WHERE {
  ?request :involvesPerson ?person .                              #5
  ?person a :Pregnant_SportsWalking_Person .                   #6
  ?person a :Person .                                           #7
  ?person :hasPreferredLanguage ?preferredLang .                #8
  ?preferredLang :hasLanguageCode ?langCode .                  #9
  ?request :involvesLocation ?location .                         #10
  ?location :hasRelatedIndex :AQI_bad .                          #11
  BIND (BNODE() AS ?personalized_rec) .                          #12
  BIND ( :getPersonalizedRecommendation_langParam(?langCode, "bad",
:Pregnant_SportsWalking_Person) AS ?rec_description) . }      #13
```

Fig. 3. An example rule that combines inferred data and function's output to create new triples.

The first level of rules (109 in total) takes into account the schema and populated instances and produces low level derivations for the next level. These rules handle the transformation of age values to relevant age groups, of AQ observations to AQI scales and the classification of user profiles into those profile categories for which recommendations are available, on the basis of provided environmental experts' knowledge. Considering the rule-driven user classification process, a user is classified in one of the available basic profiles (subclasses of class `Person`), or in any permissible combination of them. Rules handle complex profiles automatically, by "downgrading" them into those combinations for which recommendations are defined. A user may be

provided with more than one recommendation messages, with respect to its profile characteristics and the rule-based classification output.

The second level of rules (150 in total) takes into account preceding results from user- and AQI-related inferences, as input to relevant SPIN functions, which in turn generate the text of the most relevant recommendation per case from a list of experts' recommendations residing in the ontology. An example of the second level's SPIN rule is presented in Fig. 3, where a recommendation is formulated on the basis of preceding inferences: the user belongs to a specific combined category (triple #6) and the user's location was inferred to have a bad AQI (triple #11) at the time of request for DS. The user's preferred language (triples #8-9) is also considered for the final recommendation outcome, since multilingual messages are integrated in the ontology.

3.3 Communication with Ontology-based Framework

We establish an interoperable communication between the decision support KB and external modules, with the implementation of RESTful services. Complex ontological definitions and rules are hidden behind the developed web services, which handle the dynamic population of the ontology with actual data (AQ observations, user profile details, requests for DS) and the automated inference of personalized recommendations. Services were implemented in Java EE 7, with the adoption of Apache Jena framework (<https://jena.apache.org/>) for manipulating RDF graphs and the SPIN API (<http://topbraid.org/spin/api/>) for performing rule-based inference.

4 Test Cases and Inference Results

Our proposed ontology-based framework comprises the operational DS module of the hackAIR platform for generating meaningful QoL information, personalized according to the citizens' profile requirements. In this context, we demonstrate its functionality through the following indicative scenario: "*Valeria, a 32-year-old woman, pregnant with respiratory problems (asthma), queries the hackAIR application for information about existing AQ in an area where she usually goes for long walks. At the time of request, the PM_{10} values are extremely high (e.g. $150\mu\text{g}/\text{m}^3$)*". The process executed is given below:

- (1) The request is sent from the hackAIR app to the ontology-based module through the supported web service;
- (2) The user-profile details (Table 1-A) as well as the location of interest (Table 1-B) and AQ measurement (Table 1-C) are dynamically populated in the ABox according to the schema declared in the TBox;
- (3) An instance of `Request` type is additionally formulated (Table 1-D) by integrates all involved information declared in the previous step;
- (4) The recommendation module is triggered. SPIN rules are activated according to the level of suitability to the case. Results are produced dynamically (Table 1-E), moving from more general to more specific inferences;

- (5) Inferences asserted as recommendations to the user, are dynamically extracted from the ontology and pushed back to the application for visualization.

Table 1. Instances in Turtle format⁵ as populated/inferenced in our proposed framework

A. User profile

```
abox:Valeria
  rdf:type TBox:Person ;
  tbox:hasAge 32^^xsd:integer ;
  tbox:hasGender tbox:female ;
  tbox:isSensitiveTo tbox:Asthma ;
  tbox:isPregnant "true"^^xsd:boolean ;
  tbox:hasLocation abox:location_V ;
  tbox:hasPreferredActivity tbox:walking_activity ; .
```

B. Location

```
abox:location_V
  tbox:hasEnvironmentalData abox:PM10EnvData_location_V .
```

C. Environmental data

```
abox:PM10EnvData_location_V
  rdf:type tbox:AirPollutantEnvironmentalData ;
  tbox:hasEnvironmentalDataType tbox:PM10 ;
  tbox:hasNumericalValue abox:PM10Value_location_V ; .
abox:PM10Value_location_V
  rdf:type tbox:AirPollutantValue ;
  tbox:hasUnit abox:microGramsPerCubicMeter ;
  tbox:hasValueValue "150.0"^^xsd:double ; .
```

D. Request

```
abox:request_XYZ
  rdf:type tbox:Request ;
  tbox:involvesEnvironmentalData abox:PM10EnvData_location_V ;
  tbox:involvesLocation abox:location_V ;
  tbox:involvesPerson abox:Valeria ;
```

E. Inferences

```
abox:Valeria
  rdf:type tbox:AdultPerson ;
  rdf:type tbox:PregnantFemalePerson ;
  rdf:type tbox:SensitiveHealthPerson ;
  rdf:type tbox:SportsWalkingPerson ;
  rdf:type tbox:Pregnant_Sensitive_Person ;
  rdf:type tbox:Pregnant_SportsWalking_Person ; .
abox:location_V tbox:hasRelatedIndex tbox:AQI_bad ; .
abox:Valeria
  tbox:isProvidedWithRecommendation [
    rdf:type tbox:LimitExposureRecommendation ;
    tbox:hasDescription "You should go for a walk in an area with
      cleaner air."@en ;
    tbox:hasDescriptionIdentifier "walking" ; ] ;
  tbox:isProvidedWithRecommendation [
    rdf:type tbox:LimitExposureRecommendation ;
    tbox:hasDescription "Consider avoiding any intense outdoor activ-
      ity in your area! The existing air quality might be harmful
      for your health."@en ;
    tbox:hasDescriptionIdentifier "general personalized" ; ] ; .
```

⁵ Details about the Turtle format are available at: <https://www.w3.org/TR/turtle/>

Focusing on the inference results presented in Table 1-E, we may distinguish the move from *general-to-specific* classifications, as those were produced by automatically executing different types/levels of implemented SPIN rules. First, the user is categorized as `PregnantFemalePerson`, `SensitiveHealthPerson` and `SportsWalkingPerson`, conforming to relevant information provided in the user profile (Table 1-A); the user is then categorized in two combined classes (`Pregnant_Sensitive_Person` and `Pregnant_SportsWalking_Person`) which have direct association to the available recommendations provided by the system. Considering these two specific subclasses that the user belongs to, and given the fact that the air quality index in the area of interest is inferred as `AQI_bad`, the system suggests *two* different recommendations, one related to the user’s preferred activity (walking) and the other to the user’s health sensitivity (pregnancy, asthma).

In a different scenario, the rule-based inference mechanism would follow the same reasoning sequence but would trigger different rules that correspond to the semantics behind the interpreted data. For example, different user profile characteristics or preferences would lead to different user profile classifications and, thus, different recommendation messages to the users, on the basis of existing AQ condition. Representative scenarios that were created within the hackAIR project, with support from the project partners and environmental experts, are demonstrated in [23].

5 Evaluation

For the evaluation of the proposed representation and reasoning framework, we focus on the following aspects: (i) the *consistency* of the provided results, by examining if the inferred recommendations comply with those targeted to be given through the classification and reasoning process; (ii) its *performance*, in terms of elapsed time when a request is submitted to the system. Unfortunately, a direct comparison of response times between the proposed framework and alternative approaches ([5, 15]) is not feasible, since there are no benchmarks to follow; systems have different complexity, demonstrate different functionalities, input or internal processes, target different recommendation outcomes, and the implementation details that are missing block the reproduction of identical experiments within our proposed context.

For the consistency checking task, the environmental experts and pilot users of the hackAIR project performed a thorough analysis of the reasoning process by examining the recommendation inferences for each ontology-represented use case. The evaluation showed a deviation from the planned system behavior in only two cases, due to wrong intermediate classifications. The problem was fixed by correcting the corresponding rules in the ontology.

Regarding the performance evaluation of the implemented DSS, we considered 11 problem (request) descriptions, differing in size and content: simple and complex profiles, with one or more involved users per request, with different profile characteristics, etc. For all use cases, the response time of the overall DS process ranged from 1.32 to 1.89 sec. (any differentiation in times depends mainly on the complexity of the rules performed per case), with an average time of 1.61 sec. The evaluation ran on a computer with: Intel® Core™ i5-4690K, x64-based processor at 3.50GHz, with

16GB RAM memory. Overall, our ontology-based DSS was proved to be a light-weight approach, complying with the needs for fast computations and response times of the reasoning service, avoiding any redundant system complexity.

6 Advantages of the Proposed Implementation

Through the extensive use of ontologies we achieve to create an integrated, stand-alone DS framework that is *self-identified* (structures the content and dynamically computes implicit values based on knowledge stated in the schema), and *self-triggered* (fires the reasoning process automatically and appropriate rules are executed when needed). The advantages of such an implementation are described below.

Uniformity: SPIN rules are defined as RDF triples alongside the ontology model. Such an approach prevents additional programming effort (rule-based inference is operated via the integrated SPIN engine) and facilitates the adoption of the DS module and external sources through implemented RESTful web services.

Flexibility: The multifaceted proposed framework permits multilingual definition of recommendations through the initialization of a single instance of `Recommendation` type per message. This can “carry” the same text in any of the supported languages, while an appropriate SPIN rule will return to the user the inferred message on his/her preferred spoken language. Different recommendation messages for the same use case are also easily supported, in a sense that the proposed framework can select either randomly, or by validating specific weighted factors, a personalized recommendation from a “bag of messages”, all of which are defined as relevant for the examined case.

Modularity: The parameterization of implemented SPIN rules and functions allow complex schema and property declarations to be manageable in terms of maintenance, adoption and re-use; changing the body of a rule does not affect its use in other stages of the rule layer, while changing the parameters used/returned only alters its function signature. Moreover, SPIN enables the explicit definition of the rule execution order, giving priority to or isolating groups of rules under specified circumstances.

Extensibility: The hierarchical capabilities of the proposed framework facilitate the update and extension of rules and concepts supporting additional input from AQ, health or other related domain knowledge, without the need to change rules in other layers. More specifically, in our implementation, air pollutant measurements are converted via rules to corresponding AQI values; such data serve as input to subsequent levels of the reasoning process. If an additional air pollutant is to be added to the schema, rules will be enriched only for the task of converting numerical to nominal values, without any interference to AQI rules or SPIN rules interconnected to them.

Appropriateness: The combination of ontologies with SPIN rules, as an extra layer of rule-based reasoning, is perfectly suited for more advanced, rule-based inference. This sophisticated implementation provides the required expressiveness for the domain’s knowledge representation and the respective inference mechanisms which can efficiently deal with the provision of user-profile driven recommendations.

7 Conclusions and Future Work

This paper presented a novel ontology-based framework, which integrates both the *static parts* and the *dynamic processes* of a DSS in a uniformly structured and semantically enriched decision support KB. The implementation is based on a 3-layered approach which consists of the following components: (i) the ontological schema of abstract concepts and relations in the domain of interest, (ii) the structured initialization of actual data in the ontology, and (iii) the rule-based mechanism for interpreting the semantics behind the stored data and generating the targeted recommendations to the users, with respect to their user profile characteristics. The developed schema and the mechanisms integrated in the proposed DSS framework, reflect the dynamics of the AQ domain from the citizens' perspective: in every request, the system combines the status of the environment and the specific characteristics of the requesting user, so as to produce simple personalized recommendations encompassing the direct influence of severe AQ conditions in the users' health and planned activities.

The proposed implementation highlights the strategic advantages of the use of ontologies on the basis of the multifaceted needs of DSSs: the layered structure and the semantic declaration of concepts and rules can efficiently handle the heterogeneity of represented data and facilitate the modularity and extensibility of the system. Moreover, the ontology-based implementation is domain-agnostic and can be exploited in different application contexts, following a relevant configuration mechanism of the targeted user profiles and recommendations.

As future work, we aim to further evaluate the satisfiability and coverage of the recommendation system from the users' perspective. We also aim to investigate the integration of fuzzy reasoning techniques for covering more complex relations (additional parameters, cumulative effects, etc.) Finally, we will propose an Ontology Design Pattern for the formal representation of the reasoning mechanism, and exploit new rule-based approaches with the adoption of SPIN's next generation called SHACL [24].

Acknowledgments. This work is partially funded by the European Commission under the contract number H2020-688363 hackAIR.

References

1. Karatzas, K., Kukkonen, J.: Quality of life information services towards a sustainable society for the atmospheric environment. COST Action ES0602 Workshop Proc. (2009)
2. Zhu, L., Karatzas, K., Lee, L.: Urban environmental information perception and multimodal communication: the air quality example. LNAI 5398, pp. 288-299 (2009). doi: 10.1007/978-3-642-00525-1_29
3. Michels, H., Mauè, P.: Semantics for notifying events in the affecting environment. EnviroInfo Conf., pp. 501-507 (2010)
4. Karatzas, K., Endregard G., Fløisand I.: Citizen-oriented environmental information services: usage and impact modelling. 19th EnviroInfo Conf., pp. 872-878 (2005)
5. Rospocher, M., Serafini, L.: An Ontological Framework for Decision Support. JIST Conf. 2012, LNCS 7774, pp. 239-254 (2012). doi: 10.1007/978-3-642-37996-3_16

6. MOUNTZIDOU, A., et al.: Towards Air Quality Estimation Using Collected Multimodal Environmental Data. IFIN 2016, ISEM 2016, pp. 147-156 (2016). doi: 10.1007/978-3-319-50237-3_7
7. McIntosh, B.S., et al.: Environmental decision support systems (EDSS) development-challenges and best practices. *Env. Mod. and Soft. J.* 26(12), pp. 1389-1402 (2011). doi: 10.1016/j.envsoft.2011.09.009
8. Matthies, M., Giupponi, C., Ostendorf, B.: Environmental decision support systems: current issues, methods and tools. *Env. Mod. and Soft. J.* 22(2), pp. 123-127 (2007). doi: 10.1016/j.envsoft.2005.09.005
9. Raskin, R., Pan, M.: Semantic web for earth and environmental terminology (sweet). In: *Proc. of the Workshop on Semantic Web Technologies for Searching and Retrieving Scientific Data* (2003)
10. Buttigieg, P.L., Morrison, N., Smith, B., Mungall, C.J., Lewis, S.E.: The environment ontology: contextualising biological and biomedical entities. *J. of Biom. Semantics* 4(1), 43 (2013). doi: 10.1186/2041-1480-4-43
11. Oprea, M.: AIR_POLLUTION_Onto: an ontology for air pollution analysis and control. In: *Int. Conf. on Artificial Intelligence Applications and Innovations*, pp. 135-143 (2009). doi: 10.1007/978-1-4419-0221-4_17
12. Bicer, V., Tran, T., Abecker, A., Nedkov, R.: KOIOS: Utilizing semantic search for easy-access and visualization of structured environmental data. *ISWC 2011, LNCS 7032*, pp. 1-16 (2011). doi: 10.1007/978-3-642-25093-4_1
13. Ruther, M., Bandholtz, T., Logean, A.: Linked environment data for the life sciences. *3rd Int. Work. on Semantic Web Applications & Tools for the Life Sciences* (2010)
14. Kontopoulos, E., Martinopoulos, G., Lazarou, D., Bassiliades, N.: An ontology-based decision support tool for optimising domestic solar hot water system selection. *J. of Cleaner Production* 112, pp. 4636-4646 (2016)
15. Wetz, P., Trinh, T.D., Do, B.L., Anjomshoaa, A., Kiesling, E., Tjoa, A.M.: Towards an Environmental Information System for Semantic Stream Data. *28th EnviroInfo Conf.*, pp. 637-644 (2014)
16. Wanner, L. et al.: Ontology-centered environmental information delivery for personalised decision support. *Expert Systems with Applications* 42, pp. 5032-5046 (2015). doi: 10.1016/j.eswa.2015.02.048
17. Gruber, T.R.: A translation approach to portable ontology specifications. *Knowledge acquisition J.* 5(2), pp. 199-220 (1993). doi: 10.1006/knac.1993.1008
18. Turban, E., Aronson J.E., Liang T.P.: *Decision support systems and intelligent systems*. Pearson Prentice-Hall, pp. 100-136 (2005)
19. Suárez-Figueroa, M., Gómez-Pérez, A., Villazón-Terrazas, B.: How to write and use the Ontology Requirements Specification Document. *LNCS 5871*, pp. 966-982 (2009)
20. W3C OWL Working Group: *OWL 2 Web Ontology Language*. W3C Recommendation (2012). <https://www.w3.org/TR/owl2-overview/>
21. Falco, R., Gangemi, A., Peroni, S., Shotton, D., Vitali, F.: Modelling OWL ontologies with Graffoo. In *European Semantic Web Conf. (ESWC)*, pp. 320-325, Springer, Cham (2014). doi: 10.1007/978-3-319-11955-7_42
22. Knublauch, H., Hendler, J.A., Idehen, K.: SPIN - overview and motivation. W3C Member Submission (2011). <https://www.w3.org/Submission/spin-overview/>
23. hackAIR Consortium. D4.2: Semantic integration and reasoning of environmental data (2017). http://www.hackair.eu/wp-content/uploads/2016/03/d4.2-semantic_integration_and_reasoning_of_environmental_data.pdf
24. Knublauch, H., Kontokostas, D.: *Shapes Constraint Language (SHACL)*. W3C Recommendation (2017). <https://www.w3.org/TR/shacl/>