

Description Logics and Rules for Multimodal Situational Awareness in Healthcare

Georgios Meditskos, Stefanos Vrochidis, and Ioannis Kompatsiaris

Information Technologies Institute
Centre for Research and Technology - Hellas, Thessaloniki, Greece
{gmeditsk, stefanos, ikom}@iti.gr

Abstract. We present a framework for semantic situation understanding and interpretation of multimodal data using Description Logics (DL) and rules. More precisely, we use DL models to formally describe contextualised dependencies among verbal and non-verbal descriptors in multimodal natural language interfaces, while context aggregation, fusion and interpretation is supported by SPARQL rules. Both background knowledge and multimodal data, e.g. language analysis results, facial expressions and gestures recognized from multimedia streams, are captured in terms of OWL 2 ontology axioms, the de facto standard formalism of DL models on the Web, fostering reusability, adaptability and interoperability of the framework. The framework has been applied in the eminent field of healthcare, providing the models for the semantic enrichment and fusion of verbal and non-verbal descriptors in dialogue-based systems.

Keywords: multimodal data, ontologies, rules, situation awareness

1 Introduction

A key requirement in multimodal domains is the ability to integrate the different pieces of information (modalities), so as to derive high-level interpretations. More precisely, in such environments, information is typically collected from multiple sources and complementary modalities, such as from multimedia streams (e.g. using video analysis and speech recognition), lifestyle and environmental sensors [15]. Though each modality is informative on specific aspects of interest, the individual pieces of information themselves are not capable of delineating complex situations. Combined pieces of information on the other hand can plausibly describe the semantics of situations, facilitating intelligent situation awareness.

In parallel, the demand for context-aware user task support has proliferated in the recent years across a multitude of application domains, ranging from healthcare and smart spaces to transportation and energy control. A key challenge in such applications is to abstract and fuse the captured *context* in order to elicit an adequate understanding of user actions [5]. In healthcare, for example, wearable and ambient sensors, coupled with profile information and clinical knowledge can be used to improve the quality of life of care recipients and provide useful insights to clinical experts for personalized interventions and care solutions [23].

Given the inherent requirement in multimodal environments to aggregate low-level information and integrate domain knowledge, it comes as no surprise that Semantic Web technologies have been acknowledged as affording a number of highly desirable features. More precisely, the OWL 2 ontology language [10] has been extensively used to capture context elements (e.g. profiles, events, activities, locations, postures and emotions) and their pertinent relations, mapping observations and domain knowledge to class and property assertions in the Description Logics (DL) [3] theory, fostering integration of information at various levels of abstraction and completeness [14]. The generated models encapsulate formal and expressive semantics, harvesting several benefits brought by ontologies, e.g. modelling of complex logical relations, sharing information from heterogeneous sources, sound and complete reasoning engines.

This paper describes an ontology-based framework for context awareness and conversation understanding in multimodal natural language interfaces. The framework allows the semantic enrichment of verbal and non-verbal information coming from multiple devices and acquisition methods, e.g. from multimedia analysis, following a knowledge-driven methodology for observation aggregation, linking and situation interpretation. The contributions of our work can be summarized in the following:

- We alleviate the lack of inherent temporal reasoning support in DL and OWL 2 by adopting an a-temporal approach for subsumption reasoning and multimodal data fusion based on time-windows.
- We propose an iterative combination of DL reasoning and rules to enhance the reasoning capabilities of the framework.
- We use SPARQL queries (`CONSTRUCT` graph patterns) as the underlying rule language of the framework, overcoming the lack of a standard rule language that runs directly on top of RDF and OWL ontologies.
- Due to the dynamic and open nature of ontologies, the framework is modality-agnostic, in the sense that it is not tight to specific domains and data sources but it can be extended, adapted and used in a variety of situations.

We illustrate the capabilities of the framework through its integration into a dialogue-based agent for conversational assistance in healthcare. More specifically, elderly use the dialogue system (usually at home) to acquire information and suggestions related to basic care and healthcare (e.g. symptoms, treatments, etc.). A key challenge in this domain is the effective fuse of verbal and non-verbal communication modalities, e.g. deictic gestures and spoken utterances, in order to disambiguate and interpret user input during the interaction with the agent.

The rest of the paper is structured as follows: Section 2 begins with a basic background on the DL theory and OWL ontologies. It continues with a discussion on ontology-based context-aware solutions, explaining basic concepts and challenges. Section 3 describes the proposed framework, providing details on the representation and interpretation layers, as well as on the hybrid reasoning scheme and the role of SPARQL. Section 4 explicates through an example use case from the ongoing simulated evaluation of the framework in the healthcare domain and Section 5 concludes the paper and outlines next steps.

2 Background and Related Work

2.1 Description Logics

Description Logics (DL [3]) is a family of knowledge representation formalisms characterised by logically grounded semantics and well-defined reasoning tasks. The main building blocks are *concepts* (or classes), representing sets of objects, *roles* (or properties), representing relationships between objects, and *individuals* (or instances) representing specific objects. Starting from atomic concepts, arbitrary complex concepts can be described through a rich set of constructors that define the conditions of concept membership. DL provides, among others, constructs for concept inclusion ($C \sqsubseteq D$), equality ($C \equiv D$) and assertion ($C(a)$), as well as role inclusion ($R \sqsubseteq S$) and assertion ($R(a, b)$).

The semantics of a DL language is formally defined through an interpretation I that consists of a nonempty set Δ^I (the domain of interpretation) and an interpretation function \cdot^I , which assigns to every atomic concept A a set $A^I \subseteq \Delta^I$ and to every atomic role R a binary relation $R^I \subseteq \Delta^I \times \Delta^I$. Table 1 shows the syntax and semantics of some of the most common DL constructors. For example, the class of all deictic gesture observations that point to the head can be defined as $\text{PointsToHeadGesture} \equiv \text{DeicticGesture} \sqcap \exists \text{hasBodyPart}.\{\text{head}\}$.

Besides formal semantics, DL comes with a set of powerful reasoning services, for which efficient, sound and complete reasoning algorithms are available. For example, through *subsumption*, one can derive implicit taxonomic relations among concepts. *Satisfiability* and *consistency checking* are useful to determine whether a knowledge base is meaningful at all. *Instance realization* returns all concepts from the knowledge base that a given individual is an instance of.

2.2 OWL 2 Ontologies

An ontology is a set of precise descriptive statements about some part of the world (usually referred to as the domain of interest). Precise descriptions satisfy several purposes: most notably, they prevent misunderstandings in human communication and they ensure that software behaves in a uniform, predictable way and works well with other software¹.

¹ <https://www.w3.org/TR/owl2-overview/>

Table 1. Examples of concept and role constructors in DL.

Name	Syntax	Semantics
Intersection	$C \sqcap D$	$C^I \cap D^I$
Union	$C \sqcup D$	$C^I \cup D^I$
Universal Quantification	$\forall R.C$	$\{a \in \Delta^I \mid \forall b.(a, b) \in R^I \rightarrow b \in C^I\}$
Existential Quantification	$\exists R.C$	$\{a \in \Delta^I \mid \exists b.(a, b) \in R^I \wedge b \in C^I\}$

The Web Ontology language (OWL/OWL 2) [10] is a knowledge representation language widely used within the Semantic Web community for creating ontologies. The design and semantics of OWL 2 have been strongly influenced by DL². Some basic notions are: a) *axioms*, the basic statements that an OWL ontology expresses, b) *entities*, elements used to refer to real-world objects, and c) *expressions*, combinations of entities to form complex descriptions.

In principle, every OWL 2 ontology is essentially a collection of such basic “pieces of knowledge”. Statements that are made in an ontology are called axioms, and the ontology asserts that its axioms are true. However, despite the rich primitives, there are certain limitations that amount to the DL style model theory used to formalise semantics, and particularly the tree model property [16] conditioning DL decidability. For example OWL 2 can model only domains where objects are connected in a tree-like manner. In order to leverage OWL’s limited relational expressiveness, research has been devoted to the integration of OWL with rules (e.g. SWRL [12], SPIN [13]). User-defined rules on top of the ontology allow expressing richer semantic relations that lie beyond OWL’s expressive capabilities and couple ontological and rule knowledge [9].

2.3 Ontology-based Context Awareness and Fusion

Congruous with the open nature of context-awareness, where information at various levels of abstraction and completeness has to be integrated, ontologies have attracted growing interest as means for modelling and reasoning over contextual information in various domains [2][14]. For example, BeAware! [4] provides a framework for context awareness in road traffic management; [26] proposes an ontology-based framework for context-aware activity recognition in smart homes. A survey on context awareness from an IoT perspective is presented in [17], whereas challenges and opportunities in applying Semantic Web technologies in context-aware pervasive applications are discussed in [27].

A common characteristic in all cases above is the use of ontologies for domain modelling. Ontology languages, such as OWL 2, share a common understanding of the structure and semantics of information, enabling knowledge reuse and inferencing. Capitalizing on the expressivity of the models, several approaches define one or more interpretation layers in order to elicit an understanding of the situation. For example, in the domain of natural language interfaces and dialogue-based systems [24], ontologies provide the vocabulary and semantics for content disambiguation [6][7], such as WordNet³ and BabelNet⁴. Ontologies have been also used in NLP information extraction contexts for coreference resolution in textual input [22][19]. In the domain of multimodal fusion, ontologies are used to fuse multi-level contextual information [8]. For example, [18] presents a framework for coupling audio-visual cues with multimedia ontologies.

² In this paper, OWL 2 is used to refer to OWL 2 DL ontologies interpreted using the Direct Semantics [20]

³ <http://wordnet-rdf.princeton.edu/>

⁴ <http://babelnet.org/rdf/page/>

Relevant approaches are also described in [1] for various multimedia analysis tasks. SMARTKOM [25] partially uses ontologies to fuse information in multimodal dialogue systems, combining speech, gesture and facial expressions.

Similar to the aforementioned approaches, we use OWL 2 ontologies for modelling context types and their relationships in terms of DL concept class constructors. However, we argue that the constructors provided by DL, and hence by OWL 2, are sometimes inadequate to facilitate effective multimodal fusion. Certain modelling and reasoning limitations, such as the tree-model property mentioned above or the lack of temporal reasoning, render OWL 2 insufficient to address practical fusion requirements, such as the assertion of property fillers for unconnected instances, as we demonstrate in Section 4. Our framework leverages OWL 2 limited expressivity through an intelligent, multi-tier hybrid scheme of DL reasoning that follows a context-aware fusion and interpretation solution along with the use of SPARQL CONSTRUCT graph patterns [11] as the underlying rule language of the framework.

3 Semantic Fusion Framework

The aim of the Semantic Fusion Framework (SFF) is to aggregate context types and couple them with background knowledge. SFF does not impose any restriction on the modalities that can be fused, provided that the underlying ontologies support their representation. As such, SFF consists of two core tiers:

- **Representation tier:** Provides the knowledge structures needed to capture the semantics and structure of the various modalities, as well as the semantics of the domain model that drives the fusion task.
- **Interpretation tier:** Implements the fusion logic, capitalizing on OWL 2 DL reasoning and custom interpretation rules that combine the available input and generate additional inferences.

The conceptual architecture of SFF is depicted in Fig. 1. In the following sections, we further elaborate on the specifics of each tier.

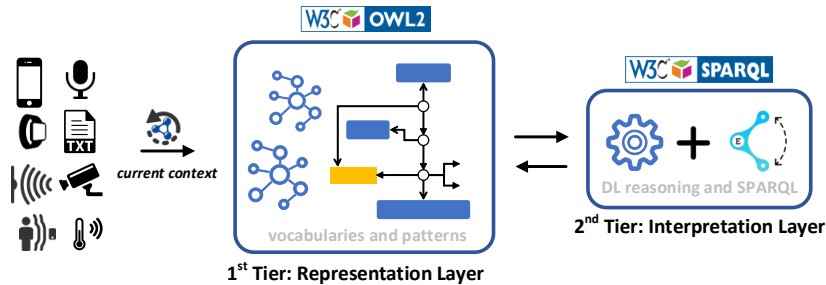


Fig. 1. Conceptual architecture of the Semantic Fusion Framework (SFF)

3.1 Domain and Context Descriptors

As mentioned above, SFF is modality-agnostic since it is not tight to specific context types. In that sense, contextual information may be collected from a variety of sources, such as ambient and wearable sensors (e.g. temperature and proximity observations), multimedia analysis, such as text analysis (named entities and concepts), video analysis (e.g. location, gestures), etc. All this information needs to be mapped on domain entities to enable the derivation of contextual descriptors that best satisfy and interpret the context.

We use the term “observation” to abstractly refer to the root of the context type hierarchy. Fig. 2 depicts a lightweight vocabulary for modelling context types. The ontology extends the `leo:Event` concept of LODE [21] to benefit from existing vocabularies to describe events and observations. Property assertions about the temporal extension of the observations and the agent (actor) are allowed, reusing core properties of LODE. Fig. 2 also depicts the relationship between the upper-level domain and context models. More precisely, the `Context` class is provided that allows one or more `contains` property assertions referring to observations. In terms of DL semantics, the `Context` class is defined as:

$$\text{Context} \equiv \exists \text{contains. Observation} \quad (1)$$

classifying instances with `contains` property assertions in `Context`. As we demonstrate in Section 4, the adaptation of the framework in different domains involves the extension of the `Context` concept, specifying the observation types that designate complex situations of interest that need to be recognized. Intuitively, instances of the `Context` concept define set of observations, designating the *current context* that needs to be classified and interpreted.

3.2 DL Reasoning and SPARQL

The interpretation tier defines the way atomic observations can lead to the derivation of high-level interpretations. For this task, we group observations into a single `Context` instance, creating the *current context*, which is then fed into the DL reasoner for subsumption reasoning and context classification. In principle, the current context is built taking into account the temporal extension of observations, along with background information pertinent to the domain. We present an example current context definition in Section 4.




Fig. 2. Upper-level domain and context structures in SFF

However, apart from context classification, an important reasoning requirement in multimodal fusion is the propagation of property fillers among incoming observations, e.g. the injection of the body part where a deictic gesture points to, which is derived after fusion with spoken utterances. Due to the tree-model property, DL reasoning is not able to update property fillers for unconnected instances (observations). SFF uses SPARQL CONSTRUCT graph patterns to enrich the reasoning capabilities of the framework, implementing certain fusion requirements, according to the entities and relations involved.

The hybrid reasoning algorithm is depicted in Figure 3. Assuming that G is the RDF/OWL graph with context observations, Q is the set with all SPARQL CONSTRUCT graph patterns, R_{DL} is the OWL 2 DL reasoning module and R_{SPARQL} is the SPARQL query engine, the algorithm in Figure 3 enriches G with additional interpretations. More specifically, the algorithm implements an iterative combination of DL reasoning and SPARQL query execution. Initially, the DL reasoning module is used over G for subsumption reasoning and realization (line 2). The derivations are added back to G that is now used as the underlying graph for the SPARQL reasoning module. When all SPARQL queries have been executed (lines 3 to 5), a reasoning iteration has been completed. The algorithm terminates when no SPARQL inferences are derived after an iteration.

4 Use Case: Reference Resolution

We describe the simulated evaluation of SFF that involves the conversation of users with a dialogue-based agent at a home or a nursing environment in order to acquire treatment suggestions about problems they have. We describe the ontologies and rules needed to disambiguate referring expressions, taking into account non-verbal modalities, e.g. deictic gestures. In the simulated example, the user touches his head and says “It hurts here!”. By fusing pointing gestures, the agent can conclude that the user has a headache and it can provide relevant treatment suggestions.

4.1 Domain Ontologies

It is assumed that SFF acquires contextual information about body gestures (e.g. deictic gestures to the head) and verbal events (e.g. entities and concepts extracted through language analysis) through respective multimedia analysis

```

1: repeat
2:    $G \leftarrow G \cup R_{DL}(G)$ 
3:   for all  $q \in Q$  do
4:      $G \leftarrow G \cup R_{SPARQL}(q, G)$ 
5:   end for
6: until  $R_{SPARQL}(q, G) = \emptyset$ 

```

Fig. 3. Skeleton of the hybrid context interpretation algorithm.

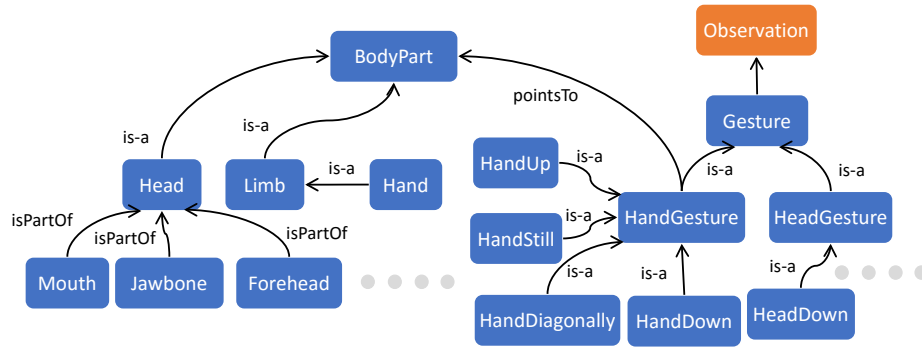


Fig. 4. Excerpt of the gesture and body part ontologies.

modules (in this case, from video and audio data). Fig. 4 depicts the specialization of the **Observation** hierarchy (Fig. 2) for modelling gestures. The emphasis is placed on the **HandGesture** concept that allows **pointsTo** property assertions about the body part where the hand points to. Additional complex concepts are defined (not visualised in Fig. 4) by composing existing contexts, e.g. $\text{HeadReference} \equiv \text{HandGesture} \sqcap \exists \text{pointsTo}.\text{Head}$. As far as language analysis is concerned, the current deployment capitalizes on the results of a frame-based formalisation of natural language utterances using DOLCE-DnS Ultralite patterns⁵. Fig. 5 depicts the relevant ontology. For example, a gesture event pointing to the head can be represented as:

```

:g1 a :HandGesture;
  :pointsTo [rdf:type :Head] ;
  :atTime [...] .

```

which is further classified as **HeadReference**, based on the axiom defined above. Likewise, the verbal event corresponding to the example can be represented as:

```

:fs1 a :InformSpeechAct, :PerceptionBodyFrameSituation;
  dul:isSettingFor :h1 .
:h1 a [:Hurt rdfs:SubClassOf dul:Event] ;
  dul:hasParticipant :d1 .
:d1 a :Here ;
  dul:isClassifiedBy [:bp1 a :BodyPart], [:sd1 a :SpatialDeictic] .

```

As illustrated, the example utterance is an **InformSpeechAct** about a physical experience (i.e. a **PerceptionBodyFrameSituation**), where the affected body part, i.e. the object classified as **BodyPart**, is not named explicitly but instead implied by a deictic referring expression.

⁵ <http://ontologydesignpatterns.org/>

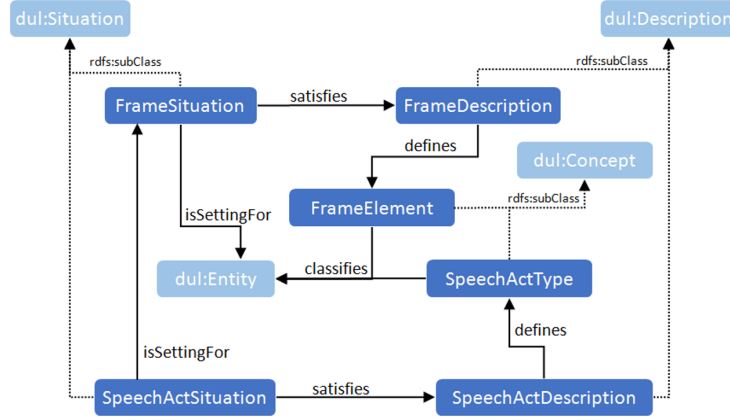


Fig. 5. The upper level ontology for representing verbal analysis results.

4.2 Context Models and Fusion

As already discussed, SFF needs to build the current context. In our example, whenever a **FrameSituation** is sent to the SFF framework, SPARQL queries retrieve neighbouring events that overlap a fixed time interval around it, e.g. $[-2s, +2s]$. The overlapped observations form the current context, which is fed into the ontology reasoner to interpret it. The current context is defined as

$$\text{CurrentContext} \equiv \text{Context} \sqcap \exists \text{contains} . \text{FrameSituation} \quad (2)$$

where **Context** is given by (1). In order to model the situation when the user feels pain, **CurrentContext** is further specialised as:

$$\text{PainContext} \equiv \text{CurrentContext} \sqcap \exists \text{contains} . (\text{FrameSituation} \sqcap \exists \text{isSettingFor} . \text{Hurt}) \quad (3)$$

According to (3), if the current context contains a **FrameSituation** that is associated with a **Hurt** conceptualisation, it is classified in the **PainContext** class. Assuming that **fs1** is part of the current context, SFF interprets it as a **PainContext** situation, since **fs1** satisfies the complex class description in (3).

In addition, provided that the **Hurt** instantiation of the **FrameSituation** is also associated with a body part, the current context can be further classified in the **Headache** class, defined as:

$$\begin{aligned} \text{Headache} \equiv & \text{PainContext} \sqcap \exists \text{contains} . \\ & (\text{FrameSituation} \sqcap \exists \text{isSettingFor} . \\ & (\text{Hurt} \sqcap \exists \text{isAssociatedWith} . \\ & (\text{Head} \sqcap \exists \text{isClassifiedBy} . \text{BodyPart}))) \end{aligned} \quad (4)$$

As such, if the user explicitly mentions the body part, then the **FrameSituation** can be directly classified by the underlying ontology reasoner as a **Headache**. In

our example, however, the user does not explicitly refer to the head, but instead points to it, while using the deictic referring expression “here”. As a result, the inferred `PainContext` is associated with a non-body part entity, which moreover is classified as `SpatialDeictic`. In this case, SFF needs to take into account the fact that there is an underspecified body part in the `FrameSituation` that requires additional contextual information, and in particular non-verbal one, in order to resolve the ambiguity and provide an appropriate feedback. The logic to derive such inferences is beyond the expressivity provided by OWL 2. In this case, SFF uses a fusion rule to resolve this ambiguity. The following SPARQL rule implements the fusion of language analysis results with hand gestures to body parts, so as to fill the missing body part fillers.

```

CONSTRUCT {
  ?p isAssociatedWith ?bodypart.
}
WHERE {
  ?c a PainContext;
    contains [isSettingFor ?p] .
  ?p a Hurt;
    isAssociatedWith ?bp .
  ?bp a SpatialDeictic.
  ?c contains [a HandGesture; pointTo ?bodypart].
}

```

Having updated the context with the inferred body part, the DL reasoner can now classify the current context in the `Headache` class, based on (4). As such, through the combination of the DL and SPARQL modules, SFF interprets the current situation as a headache, propagating it to subsequent modules to retrieve suggestions and provide feedback to the end user.

5 Conclusions

In this work, we presented SFF, an ontology-driven framework that couples DL reasoning and rules for multimodal fusion. The focus has been given on the interpretation of conversational contexts in dialogue-based systems, fusing non-verbal (i.e. gestures) and verbal features extracted from multimedia data for situation awareness. Ontologies are used to formally capture context types and background knowledge, while fusion and interpretation is reduced on the efficient combination of DL reasoning and SPARQL query execution.

We also described the simulated evaluation of SFF for reference resolution. We are currently collecting data for evaluating the framework using real-world conversations. In parallel, we are working towards further enrichment of the fusion and interpretation capabilities of the framework, so as to support for additional use cases, e.g. tasking into account emotions and facial expressions. It is also important to mention that the identification of the current conversational context does not take into account uncertainty. Our plan is to investigate

lightweight probabilistic and non-monotonic reasoning schemes to enhance the interpretation capabilities of SFF.

Acknowledgments. This work has been partially supported by the H2020-645012 project “KRISTINA: A Knowledge-Based Information Agent with Social Competence and Human Interaction Capabilities”.

References

1. Atrey, P.K., Hossain, M.A., El Saddik, A., Kankanhalli, M.S.: Multimodal fusion for multimedia analysis: A survey. *Multimedia Syst.* 16(6), 345–379 (Nov 2010)
2. Attard, J., Scerri, S., Rivera, I., Handschuh, S.: Ontology-based situation recognition for context-aware systems. In: *Proceedings of the 9th International Conference on Semantic Systems*. pp. 113–120. I-SEMANTICS '13, ACM (2013)
3. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press (2003)
4. Baumgartner, N., Gottesheim, W., Mitsch, S., Retschitzegger, W., Schwinger, W.: Beaware! - situation awareness, the ontology-driven way. *Data & Knowledge Engineering* 69(11), 1181 – 1193 (2010)
5. Bettini, C., Brdiczka, O., Henriksen, K., Indulska, J., Nicklas, D., Ranganathan, A., Riboni, D.: A survey of context modelling and reasoning techniques. *Pervasive Mob. Comput.* 6(2), 161–180 (Apr 2010)
6. Damljanović, D., Agatonović, M., Cunningham, H., Bontcheva, K.: Improving habitability of natural language interfaces for querying ontologies with feedback and clarification dialogues. *Web Semantics: Science, Services and Agents on the World Wide Web* 19, 1–21 (2013)
7. Denaux, R., Dimitrova, V., Cohn, A.G.: Interacting with ontologies and linked data through controlled natural languages and dialogues. In: *Do-Form: Enabling Domain Experts to Use Formalised Reasoning-AISB Convention 2013*. pp. 18–20. Society for the study of artificial intelligence (2013)
8. Dourlens, S., Ramdane-Cherif, A., Monacelli, E.: Multi levels semantic architecture for multimodal interaction. *Applied intelligence* 38(4), 586–599 (2013)
9. Eiter, T., Ianni, G., Krennwallner, T., Polleres, A.: Rules and ontologies for the semantic web. In: Baroglio, C., Bonatti, P., Matuszyski, J., Marchiori, M., Polleres, A., Schaffert, S. (eds.) *Reasoning Web, Lecture Notes in Computer Science*, vol. 5224, pp. 1–53. Springer Berlin Heidelberg (2008)
10. Grau, B.C., Horrocks, I., Motik, B., Parsia, B., Patel-Schneider, P., Sattler, U.: OWL 2: The Next Step for OWL. *Web Semantics: Science, Services and Agents on the World Wide Web* 6(4), 309–322 (October 2008)
11. Harris, S., Seaborne, A.: SPARQL 1.1 query language, W3C recommendation 21 march 2013, <http://www.w3.org/TR/sparql11-query/>
12. Horrocks, I., Patel-Schneider, P.F., Boley, H., Tabet, S., Grosz, B., Dean, M.: SWRL: A semantic web rule language combining OWL and RuleML. Tech. rep., National Research Council of Canada and Stanford University (May 2004)
13. Knublauch, H., Hendler, J.A., Idehen, K.: SPIN - overview and motivation. W3C member submission, World Wide Web Consortium (Feb 2011)

14. Kokar, M.M., Matheus, C.J., Baclawski, K.: Ontology-based situation awareness. *Information Fusion, Special Issue on High-level Information Fusion and Situation Awareness* 10(1), 83 – 98 (2009)
15. Lahat, D., Adali, T., Jutten, C.: Multimodal data fusion: An overview of methods, challenges, and prospects. *Proceedings of the IEEE* 103(9), 1449–1477 (Sept 2015)
16. Motik, B., Cuenca Grau, B., Sattler, U.: Structured objects in OWL: representation and reasoning. In: *Proceedings of the 17th international conference on World Wide Web (WWW '08)*. pp. 555–564. ACM, New York, NY, USA (2008)
17. Perera, C., Zaslavsky, A., Christen, P., Georgakopoulos, D.: Context aware computing for the internet of things: A survey. *IEEE Communications Surveys & Tutorials* 16(1), 414–454 (2014)
18. Perperis, T., Giannakopoulos, T., Makris, A., Kosmopoulos, D.I., Tsekeridou, S., Perantonis, S.J., Theodoridis, S.: Multimodal and ontology-based fusion approaches of audio and visual processing for violence detection in movies. *Expert Systems with Applications* 38(11), 14102 – 14116 (2011)
19. Prokofyev, R., Tonon, A., Luggen, M., Vouilloz, L., Difallah, D.E., Cudré-Mauroux, P.: Sanaphor: Ontology-based coreference resolution. In: *International Semantic Web Conference*. pp. 458–473. Springer (2015)
20. Schneider, M., Rudolph, S., Sutcliffe, G.: Modeling in OWL 2 without restrictions. In: *Proceedings of the 10th International Workshop on OWL: Experiences and Directions, co-located with 10th Extended Semantic Web Conference* (2013)
21. Shaw, R., Troncy, R., Hardman, L.: Lode: Linking open descriptions of events. In: *4th Asian Conference on The Semantic Web*. pp. 153–167. Shanghai, China (2009)
22. Sleeman, J., Finin, T.: Type prediction for efficient coreference resolution in heterogeneous semantic graphs. In: *Semantic Computing (ICSC), 2013 IEEE Seventh International Conference on*. pp. 78–85. IEEE (2013)
23. Solanas, A., Patsakis, C., Conti, M., Vlachos, I.S., Ramos, V., Falcone, F., Postolache, O., Pérez-Martínez, P.A., Di Pietro, R., Perrea, D.N., et al.: Smart health: a context-aware health paradigm within smart cities. *IEEE Communications Magazine* 52(8), 74–81 (2014)
24. Sonntag, D.: *Ontologies and Adaptivity in Dialogue for Question Answering*, vol. 4. IOS Press (2010)
25. Wahlster, W.: *Dialogue Systems Go Multimodal: The SmartKom Experience*, pp. 3–27. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
26. Wongpatikaseree, K., Ikeda, M., Buranarach, M., Supnithi, T., Lim, A.O., Tan, Y.: Activity recognition using context-aware infrastructure ontology in smart home domain. In: *Knowledge, Information and Creativity Support*. pp. 50–57 (2012)
27. Ye, J., Dasiopoulou, S., Stevenson, G., Meditskos, G., Kontopoulos, E., Kompatsiaris, I., Dobson, S.: Semantic web technologies in pervasive computing: A survey and research roadmap. *Pervasive and Mobile Computing* 23, 1–25 (2015)