

SAREF4health: Towards IoT standard-based ontology-driven cardiac e-health systems

João Moreira ^{a,*}, Luís Ferreira Pires ^a, Marten van Sinderen ^a, Laura Daniele ^b and Marc Girod-Genet ^c

^a University of Twente, Enschede, The Netherlands

E-mails: j.luizrebelomoreira@utwente.nl, l.ferreirapires@utwente.nl, m.j.vansinderen@utwente.nl

^b TNO, Netherlands Organisation for Applied Scientific Research

E-mail: laura.daniele@tno.nl

^c Institut Mines-Télécom – Télécom SudParis – Institut Polytechnique de Paris

E-mail: marc.girod_genet@telecom-sudparis.eu

Abstract. Recently, a number of ontology-driven healthcare systems have been leveraged by the Internet-of-Things (IoT) technologies that offer opportunities to improve abnormal situation detection when integrating medical wearables and cloud infrastructure. Usually, these systems rely on standardised IoT ontologies to represent sensor data observations. The ETSI Smart Applications REFerence ontology (SAREF) is an extensible industry-oriented standard. In this paper, we explain the need for interoperability of IoT healthcare applications and the role of standardised ontologies to achieve semantic interoperability. In particular, we discuss the verbosity problem of SAREF when used for real-time electrocardiography (ECG), emphasizing the requirement of representing time series. We compared the main ontologies in this context, according to quality, message size (payload), IoT-orientation and standardisation. Here we describe the first attempt to extend SAREF for specific e-Health use cases related to ECG data, the SAREF4health extension, which tackles the verbosity problem. Ontology-driven conceptual modelling was applied to develop SAREF4health, in which an ECG ontology grounded in the Unified Foundational Ontology (UFO), which plays the role of a reference model. The methodology was enhanced by following a standardisation procedure and considering the RDF implementation of the HL7 Fast Healthcare Interoperability Resources (FHIR) standard. The validation of SAREF4health includes the responses to competency questions, as well as the development and tests of an IoT Early Warning System prototype that uses ECG data and collision identification to detect accidents with truck drivers in a port area. This prototype integrates an existing ECG wearable with a cloud infrastructure, demonstrating the performance impact of SAREF4health considering IoT constraints. Our results show that SAREF4health enables the semantic interoperability of IoT solutions that need to deal with frequency-based time series. Design decisions regarding the trade-off between ontology quality and aggregation representation are also discussed.

Keywords: Ontology-driven healthcare, Internet of Things, ontology-driven conceptual modelling, SAREF, ECG

Accepted by: Stefano Borgo, Cogan Matthew Shimizu and Pascal Hitzler

1. Introduction

Recently, several healthcare solutions are leveraged by the Internet-of-Things (IoT) technologies (Cosío-León et al., 2018; da Costa et al., 2018; ETSI, 2019b; Haghi et al., 2017; Hossain and Muhammad, 2016; Nachabe Ismail et al., 2016; Rahmani et al., 2018), which provide a series of new capabilities for the development of integrated sensor systems that provide high-frequency data for real-time needs.

*Corresponding author. E-mail: j.luizrebelomoreira@utwente.nl.

In particular, healthcare IoT solutions usually make use of clinical data that are monitored by medical wearables and/or sensors. These data are processed both locally, e.g., a Smartphone playing the role of gateway or proxy (as in edge computing), and remotely, e.g., through the integration within cloud infrastructures. Monitoring cardiac-related data is a common requirement of healthcare solutions. IoT solutions that provide high-frequency electrocardiography (ECG or EKG) data use ECG devices wired with electrodes that are capable of providing real-time data to a field gateway (e.g., a mobile application) via Bluetooth. The field gateway acts as a proxy that aggregates raw data and forwards to cloud infrastructures via the Internet (Hossain and Muhammad, 2016; Rahmani et al., 2018; Xia et al., 2013), allowing patient monitoring for automatic identification of abnormal situations. These types of ECG solution can be used in various application scenarios, such as for early warnings of cardiovascular accident risks.

Interoperability of IoT healthcare systems is crucial to enable seamless integration, in particular by improving the mutual understanding of the exchanged data, i.e., the semantic interoperability between senders and receivers. In the past years, numerous ontologies were proposed for the health domain and some of these also consider the IoT domain. Usually, ontology-driven IoT healthcare systems rely on ontologies that represent IoT conceptualizations, such as sensor data observations (ETSI, 2019a; Li et al., 2016; Palavalli et al., 2016). In this context, one of the main current standardised ontologies for the IoT is the *Smart Applications REFerence ontology* (SAREF), formerly known as *Smart Appliances REFerence ontology*,¹ which is an extensible IoT reference ontology developed with a focus on the IoT industry needs (Daniele et al., 2015). SAREF has been standardised by the *European Telecommunications Standards Institute* (ETSI) and also provides standardised extensions,² for smart energy (SAREF4ener), environment (SAREF4envi) and buildings (SAREF4bldg) (EC, 2017). SAREF is considered by the *European Commission* as “a first ontology standard in the IoT ecosystem and sets a template and a base for development of similar standards for the other verticals to unlock the full potential of IoT” (EC, 2017).

The knowledge goal of this paper is threefold: (1) to explain the need for interoperability in IoT applications; (2) to explain the role of ontologies to achieve IoT interoperability; and (3) to explain the motivation, development and validation of one particular ontology, SAREF4Health, which is proposed as a SAREF extension for IoT-based healthcare systems.

The contribution of this work lies in addressing an actual problem (IoT interoperability) that limits successful real-world implementation of IoT applications, such as for healthcare, showing to what extent the proposed solution (SAREF4Health) can satisfy requirements for time series data exchange and improve on the current situation of SAREF. In particular, this paper addresses the verbosity problem of SAREF messages in IoT scenarios of real-time ECG, where data are represented as frequency-based time series of measurements observed by sensors. Our comparison study highlights the four main characteristics of semantic models in this context: quality, message size (payload), IoT-orientation and standardisation. While SAREF4health addresses the verbosity problem (message size), the ontology-driven conceptual modelling practice leverages the ontology quality, adopting an ECG ontology grounded in the *Unified Foundational Ontology* (UFO) that plays the role of reference model. The approach combines the method of ontology-driven conceptual modelling with standardisation initiatives, e.g., *HL7 Fast Healthcare Interoperability Resources* (FHIR), and best practices for RDF formalisation of stream data for IoT solutions.

¹ec.europa.eu/digital-single-market/en/blog/new-standard-smart-appliances-smart-home

²ec.europa.eu/digital-single-market/en/news/new-version-machine-2-machine-standard-smart-appliances-introduced-etsi

This paper extends previous work (Moreira et al., 2018a) by providing a wider context and a better motivation for the SAREF4Health development, improving the description and motivation of the validation setup, and also discussing the next steps for the standardisation of the SAREF for e-Health extension proposed in this paper. Regarding the validation, here we give a complete explanation of the test cases within the scenario of emergency services for traffic accidents in a port area. In the validation result analysis, we discuss the trade-offs regarding the impact of design decisions, ontology quality and possible limitations.

Section 2 presents some background on IoT interoperability and motivates this research by giving an overview of current initiatives in this topic. Section 3 emphasizes the role of ontologies to achieve IoT interoperability, describing the ETSI SAREF standardisation initiative and the main semantic models for healthcare. Section 4 describes interoperability requirements of e-health systems, discussing the main challenges, a specific study case and a comparison study. Section 5 details the design of SAREF4health and its main elements. Section 6 describes the SAREF4health validation, providing an overview of the validation setup, answers to the competency questions and the results from the implemented prototype. Section 7 discusses our findings and gives a follow-up agenda for the SAREF for e-Health standardisation. Section 8 summarizes our contributions, lessons learned and future work.

2. IoT interoperability

Several interoperability classifications exist (Rezaei et al., 2014) and, at the application level, the interoperability aspects to consider are: (1) Coding and formatting: binary encoding of the messages/streams that carry data (technological interoperability), i.e., packaging of data in a message (syntactic interoperability); (2) Interpretation: assignment of meaning to the data (semantic interoperability); and (3) Dialogue: process synchronization for the exchange of messages (process interoperability).

Coding and formatting are covered by existing standards, but interpretation is only partially covered (Rezaei et al., 2014). Syntactic interoperability refers to data format standards and communication protocols used for data exchange, such as, e.g., the OASIS Emergency Data Exchange Language (EDXL) and the HL7 FHIR. These standards enable ICT-based systems to communicate to each other (exchange data) through common data formats. However, common data standardisation is not enough to guarantee that these systems and their (direct or indirect) users share the meaning of the messages that they exchange.

For example, consider a patient tracking system that exchanges messages between an *Electronic Health Record* (EHR) system and an *Enterprise Resources Planning* (ERP) system or a *Medical Information System* (MIS) of a hospital, which provide services following the message structure standards of EDXL. The standard for tracking patients, coined EDXL-TEP, provides the property “vehicle kind” from the “transport type” entity within the reference model. If the first system generates a message with the value “car” referring to an “ambulance” and the other two systems follow a different definition for the “car” term (such as a personal car), then they will not share the same semantics. Because of this, decision making can be affected, leading to erroneous procedures. In this example, the hospital is waiting for the patient in an ambulance and, therefore, expects that the first medical procedures are already taken. This is an example of semantic interoperability problem.

Semantic interoperability refers to the study of meanings, the ability to automatically interpret shared data meaningfully and accurately according to agreed-upon semantics, i.e., a common information exchange reference model (an ontology). Semantic interoperability focuses on terminology and deals with

human interpretation in an unambiguous way, ensuring that the understanding of the information is the same for all participants, i.e., “that the requester and the provider have a common understanding of the ‘meanings’ of the requested services and data” (Heiler, 1995). Semantic interoperability enables seamless integration of different data sources and leverages risk identification.

The *Internet-of-Things* (IoT) technologies provide a series of new capabilities for the development of integrated sensor systems. These technologies enable the adoption of a huge amount of sensors that provide high-frequency data for real-time needs. However, these data can only be exploited if interoperability problems are addressed at all levels. The four main general interoperability problems for seamless integration of IoT solutions (Olivieri et al., 2016) are:

- (1) The connection problem: enable devices to connect to each other and the cloud (related to network aspects);
- (2) The understanding problem: mutual understanding of exchanged data, assuring common interpretation of the information (semantic interoperability);
- (3) The scalability problem: extend the involved resources of the IoT solution with minimal performance impact, and to maintain the quality of service (QoS) of the IoT solution;
- (4) The adaptation problem: change/reconfigure the involved systems at runtime.

These problems are interconnected and a solution to address one problem impacts the others. For example, improving the semantic interoperability of an IoT solution can result in more verbose messages caused by additional metadata annotations, which can impact the connection performance (Olivieri et al., 2016).

In general, IoT platforms address syntactical interoperability by adopting standardised serialisation formats, such as JavaScript Object Notation (JSON) and the Efficient XML Interchange (EXI) (Fernandez et al., 2019), or even more compact binary formats like IETF Concise Binary Object Representation (CBOR).³ Several EU projects applied semantic technologies to IoT (Brandt et al., 2013; Li et al., 2016; Palavalli et al., 2016; Szilagyí and Wira, 2016) and some of them targeted the semantic interoperability of IoT platforms within the *IoT European Platforms Initiative* (IoT-EPI),⁴ e.g., INTER-IoT (Moreira et al., 2018c). Furthermore, the *IoT European Research Cluster* (IERC) plays a major role for the integration of the solutions produced by the aforementioned projects, aiming at defining a common vision of these IoT technologies at the European level (Gyrard et al., 2015). Few industry-oriented solutions exist, such as the AWS IoT platform (AWS *Neptune triplestore*).

A common element of these IoT semantic projects is the use of JSON for Linked Data (JSON-LD) as data exchange syntax (message payload) (Ganzha et al., 2017b; Khodadadi and Sinnott, 2017; Szilagyí and Wira, 2016). The proposal for a *Semantic Web Stack for IoT* (Szilagyí and Wira, 2016) recommends JSON-LD instead of XML because of performance issues. The payload of a JSON-LD message varies according to the semantic models chosen and, although JSON-LD is built to be lightweight, this choice can affect the time necessary to transfer and process sensor data. Therefore, this type of syntactic constraint for IoT solutions should be covered during ontology engineering, playing a major role in semantic interoperable IoT platforms.

³<https://tools.ietf.org/html/rfc7049>

⁴<https://iot-epi.eu/>

3. Role of ontologies

Semantic-driven IoT solutions rely on IoT ontologies to represent sensor data observations. Over the past few years, numerous IoT ontologies were proposed to improve the semantic interoperability of IoT artefacts, i.e. data exchange among platforms, devices, gateways, applications and networks involved in IoT solutions (Ganzha et al., 2016). Guidelines for IoT ontology engineering were proposed by the *IERC Cluster Semantic Interoperability Best Practices and Recommendations* (IERC AC4), which were recently improved by the lessons learned from the FIESTA-IoT project, covering best practices and methodologies from the semantic web community (Gyrard et al., 2015).

Among the numerous IoT ontologies proposed in the literature, W3C SSN/SOSA and ETSI SAREF are the most prominent ones (Li et al., 2016; Palavalli et al., 2016), which were rigorously developed during many years by ontologists and domain experts, being applied in several IoT use cases and supported by standardisation initiatives. Although they overlap, the motivation to create SAREF in 2015 was to provide an easier vocabulary than SSN 1.0, aiming at industry-oriented IoT developers and stakeholders that are usually not ontology experts. However, SSN/SOSA and SAREF can be aligned with each other, as described by Moreira et al. (2017a). The *Web of Things* (WoT)⁵ initiative also constitutes a relevant contribution for the semantic interoperability of IoT, which focuses on high level semantics of IoT services, enabling cross-domain use cases in which sensors act as services and services can be represented through a common and generic reference model, namely the *Ontology Web Language semantic extension for Web-Services* (OWL-S).⁶

3.1. ETSI smart applications reference ontology (SAREF)

In 2013, the *European Commission* launched a standardisation initiative together with ETSI in close interaction with the smart appliances industry to create an interoperability language for the various smart and networked devices from different manufacturers that co-exist in our homes. This initiative resulted in the *Smart Appliances REference ontology* (SAREF)⁷ (Daniele et al., 2015), which was developed with the smart home market (Daniele et al., 2016) and subsequently published as an ETSI Technical Specification.⁸ At first, this ontology was built as a reference model targeting smart appliance solutions for the energy efficiency optimisation in the smart home domain, e.g., from lamps and consumer electronics to white goods, such as washing machines and ovens. However, SAREF has evolved to cover more vertical domains and the IoT domain in general to foster cross-domain interoperability. For example, it should be able to define that a motion sensor provides the same functionality whether it is located in a smart home or in a street in the context of smart cities/smart mobility. Therefore, this ontology has been recently renamed as *Smart Applications REference ontology* (SAREF) to stress that it is not limited to the smart appliances domain, but it rather improves the integration of semantic data from and across various applications in different domains. SAREF provides building blocks that enable re-utilization of different parts of the ontology according to specific requirements. It is grounded on 47 “semantic assets”, i.e., standards, proprietary data models, protocols and other ontologies, including SSN.

The motivation behind SAREF was that the market would continue to be fragmented and powerless without a (protocol-independent) semantic layer that could enable interoperability among the various

⁵<https://www.w3.org/WoT/>

⁶<https://www.w3.org/Submission/2004/07/>

⁷<http://ontology.tno.nl/saref/>

⁸ETSI TS 103 264 V2. 1.1 SmartM2M; Smart Appliances; Reference Ontology and oneM2M Mapping (2017).

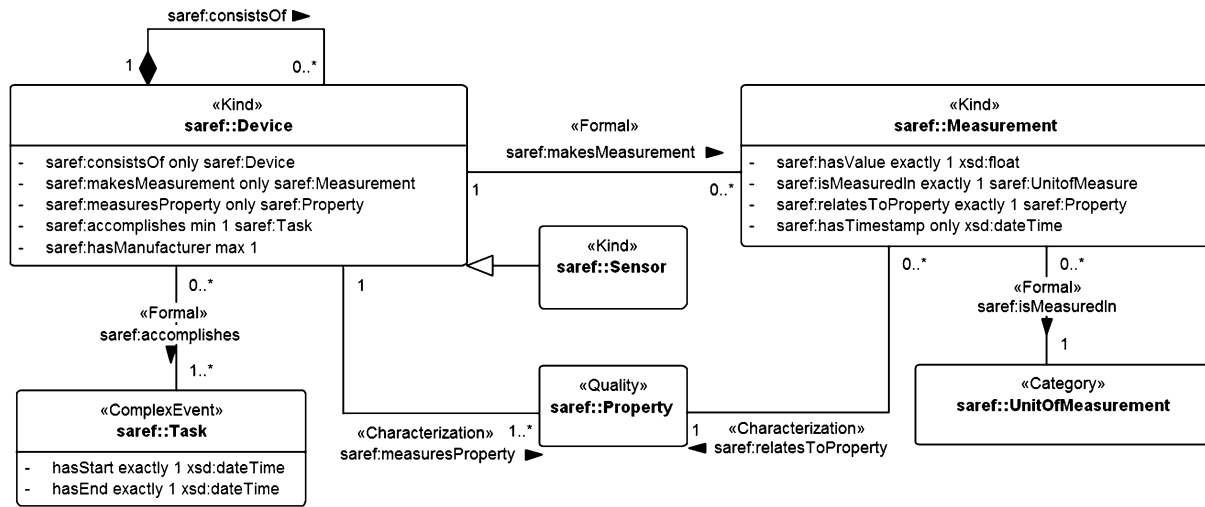


Fig. 1. Core elements for representing device measurements in SAREF.

smart appliances from different manufacturers. To that end, SAREF was created with the intention to interconnect different platforms, supporting data exchange with different protocols. Over time, SAREF has evolved with the feedback of its users, some incorporated in SAREF 2.0.⁹ Current work in ETSI includes a release of SAREF 3.0, which is expected to be published by 2020. Some of the core elements of SAREF 2.0 are depicted in Figure 1.

As described before, one of the assets used by SAREF development is SSN (version 1.0), which inspired the definition of the main elements of SAREF, namely *Device*, *Sensor*, *Unit of Measure* and *Time/Duration* (see Figure 1), according to the high-level mappings provided in the SAREF initial documentation (Daniele et al., 2015). A *saref:Device* (e.g., a *saref:Sensor*) represents a tangible object designed to accomplish one or more functions in diverse types of locations (e.g., households and buildings). For example, a *saref:Sensor* *saref:hasFunction* *saref:Function* of type *saref:SensingFunction*. The SAREF ontology offers a list of basic functions that can be combined to define more complex functions and assign them to a single device. For example, a *saref:Switch* can offer an *saref:Actuating* function of type “switching on/off” and a *saref:SensingFunction* of type *saref:LightSensor*, so if there is illumination in the environment then the switch turns off the light. Each *saref:Function* has some associated *saref:Command(s)*, which can also be selected as building blocks from a list. For example, “switching on/off” is associated with the commands “switch on”, “switch off” and “toggle”. Depending on the *saref:Function(s)* it accomplishes, a device can be found in some of the corresponding *saref:State* that is also listed as building block.

The composition (mereological) pattern of a *saref:Device* is represented through the *saref:consistsOf* self-relationship with the axiom *saref:consistsOf only saref:Device*. The *saref:consistsOf* element is defined as the “relationship indicating a composite entity that consists of other entities”, which is similar to *ufo:componentOf*, i.e., a meronymic relation that conceptualises part-whole. For example, the WM30 wind sensor (a *saref:Device*) can be defined as a composition of wind direction and wind speed sensors. A *saref:Device* measures a specific property, represented by the object property *saref:measuresProperty* to a *saref:Property*. For example, a *saref:SmokeSensor* (a *saref:Sensor*) mea-

⁹ETSI TR 103 411: SmartM2M Smart Appliances SAREF extension investigation (2017).

measures *saref:Smoke* (*saref:Property*). Analogously, a *saref:WindSensor* measures *saref:Wind*. SAREF represents a measurement observed by a sensor in time through the *saref:makesMeasurement* object property of a *saref:Device* to *saref:Measurement*, representing the relation between a device and the measurements it makes, ontologically equivalent to *sosa:Observation*. A *saref:Measurement* element is related to its *saref:UnitOfMeasure* and the *saref:Property* measured.

3.2. SAREF extensions

One of the main benefits of SAREF is its extensibility for vertical markets, which is enforced by an ETSI standard procedure for extensions (EC, 2017). A number of SAREF extensions have been created over the years, starting from the Energy, Environment and Building domains¹⁰ and resulting in the recently published extensions for the Smart Cities, Industry & Manufacturing and AgriFood domains. To avoid semantic interoperability issues when extending SAREF, ETSI defined an approach to be followed to create and submit a SAREF extension for acceptance. This procedure is based on the following best practices: (1) the extension is designed according to clarity, coherence, extensibility, minimal encoding bias and minimal ontological commitment criteria; (2) relevant stakeholders in the domain of interest should be involved in its development process; (3) the group/community that creates the extension should be committed to contribute to its maintenance; (4) it should not add concepts that are already present in SAREF or other extensions; (5) it needs to be properly documented and published.

Among the various extensions, SAREF4envi¹¹ is an extension to the environment domain and light pollution developed in collaboration with domain experts from the STARS4ALL H2020 project.¹² SAREF4envi introduced frequency-related classes and attributes like *saref4envi:FrequencyMeasurement*, so that a Device (*saref:Device*) has a frequency measurement (*saref4envi:hasFrequencyMeasurement* relation), which is measured in (*saref:isMeasuredIn* relation) *saref4envi:FrequencyUnit* unit, a SAREF unit of measure (*saref:UnitOfMeasure* type) that can be expressed in *Hertz* or reciprocal (second, hour, day, year).

Current ETSI efforts include the development of SAREF extensions for the automotive (SAREF4auto), e-Health and aging-well (SAREF4ehaw), wearables (SAREF4wear) and water (SAREF4watr) domains. These extensions are developed in the context of an EC-funded Specialist Task Force, called STF 556¹³ that has started in December 2018 and will run for a period of two and a half years.

3.3. Semantic models for healthcare systems

A well-founded ECG ontology (UFO ECG) was designed based on an ontological analysis of existing health standards and ontology-driven conceptual modelling with UFO (Gonçalves et al., 2009). The main goal of UFO ECG is to serve as a reference “unified Electronic Health Record (EHR) model”, providing mappings to the three most common standards that support the representation of ECG data, including the HL7 annotated ECG standard (HL7 aECG) (Gonçalves et al., 2011). The use of HL7 v3 along with IEEE1451 in embedded devices to achieve end to end semantic interoperability on health systems is exploited by (Cosío-León et al., 2018). HL7 aECG (HL7, 2005) is one of these standards, which was chosen by the US Food and Drug Administration (FDA) for clinical trials, implemented as a

¹⁰ETSI TS 103 410 V1.1.1: SmartM2M; Extension to SAREF; Parts 4-6 (2019).

¹¹https://www.etsi.org/deliver/etsi_ts/103400_103499/10341002/01.01.01_60/ts_10341002v010101p.pdf

¹²<http://stars4all.eu/>

¹³<https://portal.etsi.org/STF/STFs/STFHomePages/STF566>

common lexicon approach, i.e., using XML schemas, running nowadays in several hospital information systems. A review of ECG storage formats is given by Bond et al. (2011), which also includes SCP-ECG and DICOM. This review concludes that aECG inherits the verbosity of XML, since it produces large messages (25 times larger than a compressed approach), so that it may only be used in clinical drug trials and presents conceptual design issues that hinder its application in real life situations.

HL7 *Fast Healthcare Interoperability Resources* (FHIR) is an emerging open standard that is considered as the evolution of the HL7 standards for the new generation of health data exchange. HL7 FHIR defines a collection of “resources” (data model elements) that can be mixed and adapted for particular clinical contexts through the “profiling” process. Together with W3C semantic web community, the HL7 community developed an RDF representation of the FHIR data model, which can be described and validated with *Shape Expressions* (ShEx) (Solbrig et al., 2017).

The central FHIR element is the *Observation* resource,¹⁴ which is used to represent diagnosis and monitor progress, and is recommended for monitoring vital signs and device measurements. In particular, the “EKG example using Sampled Data”¹⁵ demonstrates how to serialise ECG data using the *Observation* resource, which offers the property *valueSampledData*, capable of representing the ECG sample sequence as a series of measurements of the heart electrical activity through the *SampledData* element, similarly to HL7 aECG. A *SampledData* provides a concise way to handle the data produced by devices that sample a particular physical state at a high frequency, typically used to represent the output of an ECG device. The data type includes a series of raw decimal values, which are mostly simple integers, along with adjustments for scale and factor.

The Open Biological and Biomedical Ontology (OBO) Foundry is a popular ontology-driven initiative in the biomedical informatics field, which provides a number of domain ontologies leveraged by the Basic Formal Ontology (BFO), a foundational ontology. For example, OBO provides the Eagle-I Resource Ontology, which includes several terms for health-related instruments and protocols, such as ECG monitoring device.¹⁶ SNOMED CT¹⁷ and LOINC¹⁸ are other relevant ontology-based approaches for clinical health terminologies, since they define terms like ECG voltage and ECG abnormal finding.

Finally, ETSI SmartBAN TC¹⁹ has proposed and standardised a reference model and associated modular ontology for *Body Area Networks* (BANs) and medical devices (sensors, actuators, wearables) interoperability management at all levels (data/informational/semantic level, device/technical level and network level) (ETSI, 2019a; 2019b). At technical/informational/semantic levels, the proposed solution relies on the specification and formalisation of both: (1) a BAN devices/data (measurements and control/monitoring data included) unified semantic and open reference model, associated with corresponding unified metadata and reference modular ontology, which unfortunately does not consider the time series concept; and (2) a BAN reference service model and associated service ontology. This service model, coupled with a WoT strategy, in particular handles semantic interoperability, as device discovery and composition features, and cross domain use cases. At network level, the proposed solution relies on the specification of generic Multi-Agent/oneM2M based enablers integrated into a global IoT reference architecture dedicated for secure interaction and access to BAN data and entities. This reference architecture is specified on top of the two aforementioned BAN reference models and has been clinically

¹⁴www.hl7.org/fhir/observation.html

¹⁵www.hl7.org/fhir/observation-example-sample-data.json.html

¹⁶http://purl.obolibrary.org/obo/ERO_0000835

¹⁷<http://bioportal.bioontology.org/ontologies/SNOMEDCT>

¹⁸<http://bioportal.bioontology.org/ontologies/LOINC>

¹⁹<https://www.etsi.org/committee/1413-smartban>

tested for an ‘elderly at home support and monitoring’ use case in the context of the CareWare ITEA3 project.²⁰

4. Interoperability requirements of e-health systems

4.1. General challenges

We identified that ontology-driven IoT healthcare systems usually require conceptualizations to be matched and multiple domain ontologies to be combined, because of the cross-disciplinary nature of the healthcare field, in order to provide a higher level of interoperability for systems that perform data fusion from multiple sources. *Semantic Gateways* address this issue by enabling the execution of semantic translations that can tackle specific semantic integration problems, such as semantic overload, ambiguity and information distortion (Kim et al., 2012).

Semantic translations depend on mappings between two semantic models, which guide how the input data represented with the source ontology can be transformed to data represented with the target ontology, similarly to metamodeling-based transformations in *Model-Driven Engineering* (MDE). For example, the specification of bi-directional transformations between EDXL-TEP and HL7 v2²¹ provides a set of mappings regarding the patient concept, albeit EDXL-TEP and HL7 v2 are not semantic models. SSN ontology and several other ontologies in the context of IoT, e.g., ETSI SAREF, provide mappings to OGC SWE, especially SensorML and O&M, enabling higher interoperability among these standards. In the past decade, new interoperability standards for sensor systems have been developed, including lexicon and semantic models, as well as proprietary data models, resulting in a number of different data representation assets. Some projects incorporated these standardised ontologies, but few of them addressed their alignment to support sensor data integration. To the best of our knowledge, approaches based on *Semantic Gateways* that support the whole development lifecycle of semantic translations, i.e., specification and implementation, considering their configuration at runtime, have not been exploited in any IoT healthcare system.

Our ambition has been to improve the semantic interoperability of IoT healthcare systems by considering not only IoT ontologies, but also addressing cross-domain relationships through the extension of a standardised IoT ontology with alignments to multiple domain-specific ontologies. Therefore, our research goal has been to improve the semantic integration of components of an IoT healthcare system, enabling seamless integration with other systems. We identified the following challenges to achieve this goal, along with the corresponding requirements:

1. *Semantic integration of a variety of data sources*: how to make systems understand each other, i.e., avoid semantic errors and distortions, when multiple ontologies, semantic models, standards and data models from different and overlapping domains are involved, considering their syntactic and semantic alignments?

(*Req1*) Enable the development of an IoT healthcare core context model as a well-founded ontology, taking into account existing data representations and possible restrictions from the system requirements.

²⁰<https://itea3.org/project/careware.html>

²¹<http://docs.oasis-open.org/emergency/TEP-HL7v2-transforms/v1.0/TEP-HL7v2-transforms-v1.0.html>

(Req2) Enable the development of data acquisition components to be able to process data from different sensor types/systems described with different semantic models, i.e., provide syntactic and informational/semantic interoperability.

(Req3) Enable a data fusion mechanism, allowing it to combine incoming real-time data with historical data, both for situation identification and for upstream information dissemination.

2. *Processing in time- and safety-critical applications*: how to achieve appropriate performance and scalability for real-time upstream data acquisition, risk detection and message brokering, in terms of total transaction time and quality of service?

(Req4) Enable the system to provide adequate performance, in terms of total processing time, which depends on the specific restrictions. If the requirements explicitly restrict the minimum time thresholds, the system should be able to process within this threshold.

(Req5) Enable the system to be scalable according to the data input volume and velocity, i.e., be able to allocate resources according to the number of data sources, their input frequency and message size.

(Req6) Enable a dynamic and adaptive mechanism to modify the system at runtime, minimizing the effect on the running instances.

3. *Data analysis for effective responses*: How to provide high quality situation awareness, i.e., perception, comprehension and projection, and decision making to improve decision support (edge level included)?

(Req7) Offer modelling capabilities for the representation of temporal relations among different types of events and situations. In particular, it must enable the system to use interval relationships, such as Allen's operators, to relate events for situation identification.

(Req8) Enable the representation of complex rules over context data, such as multivariate functions (risk, odds, rate and prevalence) and temporal existential rules (sliding time windows).

(Req9) Enable the system to issue warnings and alarms, as semantic enriched messages, to multiple targets through multiple channels (e.g., broker, e-mail, SMS). In particular, enable the modification of the information requirements of the targets at runtime through an UI.

4.2. Study case requirements and use cases

The INTER-IoT project defined a scenario that aims at decreasing the risk of fatal accidents at the port of Valencia, improving health prevention and enabling quick reaction by reducing response time. This scenario requires integration between logistics and e-Health domains. The goal was to exploit how cardiac e-Health can use IoT platforms dedicated to logistics to prevent the occurrence of accidents and to support evacuation or attention in case of emergency situations (Moreira et al., 2018c).

The main actor involved in this scenario is the truck driver, who works for a haulier company providing transportation of goods to/from the port. This company has an IoT platform that tracks the trips made by the trucks and monitors the health of the drivers. The solution addressing the requirements must be a third-party application for monitoring the data provided by the haulier IoT platform, integrating logistics and health data, providing the emergency notifications to be consumed by the port emergency command centre. This emergency centre is simulated in two ways: (1) an User Interface (UI) based on *Google Maps*, able to plot the data in real-time; (2) a third-party partner provides an emergency management system that consumes data from the port IoT platform and integrates them with the emergency notifications, being able to coordinate the emergency response.

INTER-IoT-EWS (Moreira et al., 2018c) is an IoT-based Early Warning System (EWS) developed to detect emergencies and risks of accidents with truck drivers wearing the *Shimmer3* ECG unit²² while they are transporting goods in a port area. Non-functional requirements include the semantic integration of IoT assets, e.g., devices, gateways, brokers and applications. Functional requirements include the real-time detection of arrhythmia (bradycardia and tachycardia) from ECG data provided by the ECG device and the detection of vehicle collision/impact by processing the cross-axial function ($x^2 + y^2 + z^2$) of instantaneous tri-axial acceleration data, similarly to *Fall-MobileGuard* (Fortino and Gravina, 2015). Functional and non-functional capabilities are:

(IIOT-FC1) IoT platforms should be able to coordinate with emergency systems by detecting accidents and risks with trucks within the port area. The EWS should be able to identify vehicle collisions and severe changes of the driver's cardiac behaviour, alerting their urgency and severity to multiple targets. The acceptance criterion is to check if the EWS, built on top of IoT platform(s), is able to coordinate with emergency systems through emergency interoperability standard(s).

(IIOT-FC2) The haulier IoT platform and the port IoT platform should be able to share cardiac information about the driver, monitored in real-time through an electrocardiography (ECG) device. The solution should be able to provide both raw and calculated data, e.g., ECG sequence (time series) and heart rate (HR). These data need to be integrated in a way that the port emergency control system can consume them. The acceptance criterion is to check if the EWS is able to process health data at the application level to identify cardiac issues and send data to emergency systems based on emergency interoperability standard(s).

(IIOT-NFC1) IoT platforms should be semantically and syntactically interoperable. The solution should be able to integrate the involved IoT platform(s) so that their data syntax and semantics are understandable, i.e., can achieve common understanding among the participating parts. The acceptance criterion is to check the use of a mechanism to translate data syntax (lexicon) and semantics of messages exchanged.

(IIOT-NFC2) E-Health and logistics should be integrated at the INTER-IoT application and semantics (A&S) level, including primitives for the interpretation of medical and logistics data. The acceptance criterion is to check the use of semantic models to represent e-health and logistics data within the IoT use cases.

(IIOT-NFC3) The energy consumption (battery level) of the devices being used for the situation identification mechanism should be monitored. The acceptance criterion is to check if the solution is able to provide energy consumption data for the application level, to be consumed by the EWS.

4.3. Comparison of semantic models for IoT-based ECG monitoring systems

Allowing the exchange of lightweight messages among medical sensors and/or wearables, gateways and cloud infrastructure is an important requirement of IoT solutions when processing big data. One of the main problems is that the verbosity of messages impacts the data exchange performance and cloud infrastructure costs. For example, the costs of the *Microsoft Azure* cloud gateway (IoT Hub) vary according to the message payload ("message meter size").²³ A common solution to this problem is to aggregate measurement data at the gateway level according to a certain frequency, transmitting time series from the local gateway to the cloud from time to time (Fortino et al., 2014), as illustrated in

²²<http://www.shimmersensing.com/products/ecg-development-kit>

²³docs.microsoft.com/en-us/azure/iot-hub/iot-hub-devguide-messages-construct



Fig. 2. Common approach to deal with high-frequency data in IoT applications.



Fig. 3. Representation of a time series measurement with FHIR (left) and with SAREF (right).

Figure 2. The drawback of this approach is that when aggregating a series of measurements of one element, metadata about each specific measurement is lost, affecting the ontological expressiveness of the messages. Figure 3 illustrates this problem, showing how the ECG data represented with FHIR serialised as an array of numbers lose the details of each specific measurement, such as the timestamp, the unit of measurement and the property that the measurement is related to. Figure 3 also shows the representation of the aforementioned aspects of one measurement with SAREF.

UFO ECG represents series of measurements with the *Sample sequence* element, which is equivalent to the FHIR *Sampled Data* element. According to UFO ECG, *Observation series* represents the complex event composed by (a *part-of* relationship) *Observations* that are evenly spaced in time, which are carried out in a *Recording session*. A *Sample sequence* is a collective, i.e., an ultimate sortal, which has parts that play the same functional role in the whole, represented as an ordered sequence of data units that results from an *Observation series*. OpenEHR emphasizes this requirement of time series representation and includes the concept of “a history of events, i.e., as a time series, allowing all software to access data in a uniform way”. The root object (*History* class) provides a list of events as attribute.

Like ETSI SmartBAN, SAREF does not define a term for the time series concept, only allowing the representation of each granular measurement of a time series through the *Measurement* element. Although this approach improves the expressiveness of the data representation, thus the ontology quality, it produces an overhead to the message size when serialised. Therefore, a verbose message is generated when instantiating ECG data with SAREF, compared to the very same data represented with FHIR, UFO ECG or OpenEHR. For example, when comparing a JSON-LD message that instantiates SAREF individual *Measurements* against the similar message in FHIR RDF, but with *SampledData* element, our experiments on ECG data serialisation show that the SAREF message size is fifty times bigger (more verbose) than the FHIR message. This was calculated based on our performance evaluation (Moreira et al., 2018b), where an ECG device configured for 256 Hz frequency provides data to a semantic gateway, which accumulates data each 5 seconds before sending the message to the cloud, thus 1280 (5×256) measurements per message. While the equivalent SAREF message size has around 5 Mb,

Table 1
Comparison of semantic models for IoT-based ECG monitoring systems

Characteristic	UFO ECG	FHIR RDF	SAREF	SmartBAN	SSN/SOSA
<i>Quality</i>	+	-	+	+	+
<i>Message size (payload)</i>	+	+	-	-	-
<i>IoT-oriented</i>	-	+/-	+	+	+
<i>Standardised</i>	-	HL7	ETSI	ETSI	W3C

the FHIR message size has around 100 Kb. Even worse, the size difference between SAREF and FHIR messages grows exponentially with the number of measurements.

While in FHIR (or UFO ECG or OpenEHR) a measurement is only a number added to the *data* element (property of *SampledData*) representing data of the time series, in SAREF (as well as in ETSI SmartBAN) a *Measurement* is a data structure that implements a number of properties (e.g., *isMeasuredIn* and *relatesToProperty*). Similarly to SAREF, W3C SSN/SOSA does this through the *Observation* class, alleviating this issue by introducing the *hasSimpleResult* property, which enables the direct link to a number (literal), but still requires the representation of other properties (e.g., *observedProperty*). Therefore, due to this verbosity issue, we conclude that default SAREF is not suitable for exchanging IoT-based ECG time series data in real-time.

We studied alternatives, identifying their benefits and drawbacks. Table 1 compares the semantic models²⁴ considered in this study. The quality of the semantic models is the first characteristic we compared, following the best practices of ontology engineering described by Gyrard et al. (2015), more specifically the best practices #8 and #12, which regard common ontology pitfalls and ontology reuse, respectively. While UFO ECG, SAREF, SmartBAN and SSN/SOSA are built upon common conceptualization mechanisms, merging different concepts in the same class/property and reusing existing ontologies, FHIR RDF presents poor quality because it is a (semi) automatic serialisation of FHIR lexicon standard. Thus, FHIR RDF is a straightforward serialisation of the FHIR standard data model, which causes an overload of object properties, e.g., it has more than ten properties that represent the “description” concept, and lacks reuse of common terms, e.g., instead of having more than ten “description” terms, it could reuse *dc:description*.

The message size (payload) indicates whether the semantic model allows the representation of time series, as in UFO ECG and FHIR RDF. We also evaluated whether the semantic model is appropriate to the IoT context, i.e., whether it provides grammatical constructs for the main concepts of the IoT domain, namely sensor, actuator, protocol, observation, unit of measurement and measurement property. Specific characteristics of the ECG domain were considered in this analysis, such as the mereology of an ECG device, the elements participating in an ECG recording session, the ECG leads responsible for measuring frequency-based sample sequences (time series). Since our goal is the development of an EWS for the detection of risks of truck accidents, we also considered how to represent and retrieve multiple sensors data, such as acceleration data to calculate tri-axial function (impact) and ECG data to calculate ECG waveform features. As mentioned earlier, SSN/SOSA and SAREF are the most appropriate IoT reference ontologies, while some FHIR RDF resources can be used to represent some of the IoT constructs. Finally, we evaluated whether the semantic model is standardised by a standardisation body, such as HL7, ETSI and W3C for FHIR, SmartBAN/SAREF and SSN/SOSA (respectively).

²⁴We avoid the term “ontology” because we understand that an ontology describes the common sense knowledge domain rather than being some RDF serialisation.

Table 2
Competency questions to be responded by SAREF4health

ID	Textual description
CQ01	What is an ECG device and how it is composed (mereology)?
CQ02	What are the elements participating in an ECG recording session?
CQ03	What is an ECG lead, what are the types of ECG leads, what type of property an ECG lead measures and what type of measurement an ECG lead can measure?
CQ04	What is an ECG sample sequence?
CQ05	What is a time series of measurements?
CQ06	What is frequency (rate) measurement of an ECG sample sequence?
CQ07	How to represent tri-axial acceleration data from accelerometers of an ECG device?
CQ08	How to integrate measurements from multiple sensors (e.g., ECG leads, accelerometer and battery monitor) of an ECG device for near real-time (frequency-based) monitoring?

In our comparison study, we concluded that none of the available standardised IoT ontologies provide adequate balance of quality and payload representing time series for ECG data. Therefore, the problem addressed in this paper is how to achieve these four characteristics in a single ontology, which led us to the development of SAREF4health as a SAREF extension.

The main information requirements of this ontology regarding ECG and IoT domains can be described as competency questions. While UFO ECG provides a high quality and deeper ontological analysis of the ECG domain, with rich descriptions following foundational categories, SAREF can be used as our IoT reference model of choice. Therefore, the set of competency questions shown in Table 2 has been defined both according to the main elements of the aforementioned reference ontologies and considering the emergency use case of the INTER-IoT project.

5. SAREF4Health design

SAREF4Health was designed in the scope of a broader development, namely the “SEmantic Model-driven development for IoT Interoperability of emergenCy serviceS” (SEMioTICS) framework (Moreira et al., 2015a; 2015b; Moreira et al., 2018c). SEMioTICS aims at improving the semantic interoperability of Early Warning Systems (EWS) and their components, i.e., it can be used to develop semantic interoperable IoT-based EWS for emergency notification services. SAREF4Health has been designed in order to give proper semantic grounding to the interoperability of EWS for what concerns health-related information.

5.1. Design methodology

SEMioTICS follows a semantic Model-Driven Engineering (MDE) approach, giving emphasis to the semantic improvement at the conceptual modelling level and MDE transformations. The SEMioTICS MDE methodology prescribes specification and implementation phases at design-time, considering the EWS deployment at runtime. The design-time level is organized in three interrelated parts guided by the model-driven approach:

- 1) *Conceptual*: presents the definition of real-world constructs in a foundational ontology, which are reflected in the modelling languages. Temporal and structural aspects are addressed by extending the Unified Foundational Ontology (UFO).

- 2) *Specification*: covers the EWS design, adopting the improved modelling languages as graphical modelling languages for context, situation and reaction.
- 3) *Implementation*: realizes the specification as executable pieces of software with IoT platforms' components.

The conceptual part supports the modelling language constructs of the specification part, which is mapped onto the appropriate technology implementations by our MDE process, where implementation is (partially) generated by MDE transformations. Horizontal (same abstraction level) exogenous (different languages) transformations are applied at the specification level to integrate models, while vertical transformations are applied to generate code, as described by Moreira et al. (2015a; 2015b). For example, SAREF4health was specified (designed) with support of OntoUML, along with the model validation approach, while the RDF implementation was (partially) generated from pre-defined MDE transformations from OntoUML to RDF (Moreira et al., 2016).

SAREF4health plays the role of core context model for the envisioned use cases, being developed as a well-founded core ontology. A well-founded core ontology provides a precise definition of a specific field described with an ontological language, being independent of a specific application (Scherp et al., 2011). OntoUML is an ontological language that provides structural and temporal predicates. In practice, OntoUML is an extension of the UML metamodel that uses stereotypes to represent UFO concepts. In order to design the SAREF4health extension, we applied the cyclic semantic enhancement process prescribed in the UFO research, increasing the quality of the core ontology through model assessment based on formal lightweight verification and validation activities (Moreira et al., 2016). The UFO ECG ontology played the role of a unified well-founded EHR reference model to represent the ECG domain. A benefit of this approach is that UFO ECG provides links to HL7 aECG, which FHIR is based upon, thus enabling straightforward alignments to FHIR RDF.

The SAREF4health design was leveraged by the same interactive and iterative approach experienced in the other standardised extensions of SAREF,²⁵ which was created in a transparent manner to allow stakeholders to provide input and follow the evolution of the work. The first step comprised the requirement collection to guide the implementation and validation of the ontology. Numerous information sources were analysed, as specifications, datasets, standards, APIs and data formats, as well as domain expert opinions and existing initiatives in the healthcare domain. The second step comprised the use case collection, specified in natural language. The third step comprised the purpose and scope definition of the ontology for the specific use cases regarding the monitoring of the cardiac behaviour of drivers, reflected in the competency questions in Table 2, in a way that abnormal situations, e.g., arrhythmia (bradycardia and tachycardia), could be detected by the emergency system.

5.2. ECG time series

The main definitions represented by the SAREF4health elements are described in Table 3. Although UFO ECG and the FHIR standard have been carefully designed in collaboration with healthcare experts, we decided to review the terminology used there. For example, we argued why UFO ECG uses the term *Sample Sequence*, while FHIR uses *Sampled Data* to describe the very same concept. Either way, the terms reflect the *Time Series* collective concept, i.e., a sequence of data units in successive equally spaced points in time, so that each sampled data unit plays the same role in the series (the 'whole').

²⁵ETSI TS 103 410 V1.1.1: SmartM2M; Extension to SAREF; Parts 1-3 (2017).

Table 3
Definitions related to time series in UFO ECG and related standards

Term	Source(s)	Textual definition
<i>Sample sequence</i>	UFO ECG	Collective: “ordered sequence of samples resulting from an <i>Observation series</i> ” (ecgOnto:095).
<i>Observation series</i>	UFO ECG	Complex event: “Series of observations evenly spaced in time carried out in an ECG Recording session” (ecgOnto:093).
<i>Sampled data</i> (<i>Observation.component.valueSampledData</i>)	HL7 FHIR	“Data that come from a series of measurements taken by a device, which may have upper and lower limits”.
<i>Time Series Observation</i>	O&M (ISO 19156)	“observation whose result is a time-series”.
<i>Series</i>	HL7 aECG	“Contains one or more sequence sets sharing a common frame of reference”.
<i>Series</i> General Series Module	DICOM	A property of General ECG that “specifies the attributes that identify and describe general information about the <i>Series</i> within a <i>Study</i> ”. A <i>Series</i> is as a sequence of data elements sharing a common frame of reference.
<i>History</i>	OpenEHR	“Root object of a linear history, i.e., time series structure. For a periodic series of events, the period will be set, and the time of each Event in the History must correspond”
<i>Event</i>	OpenEHR	“Defines the abstract notion of a single event in a series. This class is generic, allowing types to be generated which are locked to particular spatial types”
<i>Recording device</i>	UFO ECG	Kind: “Device used to acquire (to record) an ECG from a given Patient by means of electrodes. Also called electrocardiograph” (ecgOnto:087).
<i>Recording device as recorder</i>	UFO ECG	Role: “Recording device as it plays the role of an ECG recorder” (ecgOnto:088).
<i>Recording session</i>	UFO ECG	Complex event: “Medical service in which the Patient is subject of ECG recording by some <i>Recording device</i> . The <i>Recording session</i> (event) can be said to temporally coincide, albeit in a different level of abstraction, with the <i>Observation series</i> (event). In other words, these two events have the same time boundaries”.
<i>Lead</i>	UFO ECG	Kind: “Viewpoint of the heart activity that emerges from an <i>Observation series</i> of the p.d. between two electrode placements on specific regions of the surface of the patient’s body” (ecgOnto:096)
<i>Lead</i>	HL7 aECG	“A vector along which the heart’s electrical activity is recorded as a waveform”
<i>Patient</i>	UFO ECG	Role: “Person who plays the role of being subject of care, i.e., scheduled to receive, receiving, or having received a healthcare service (based on ISO/TC 18308:2003)”

In SAREF4health, we introduced the term *Time Series Measurements* to refer to a time series of a sequence of measurements made by a device, in line with the terminology often used in the measurement science (*metrology*).²⁶ We did not assign this term to the ECG context by prefixing it with *ECG* because it can be applied to other types of measurements. An *ECG Sample Sequence* is a *Time Series Measurements* that is measured in *Electric Potential* units (an array) and relates to the *Heart Electrical Activity* property.

²⁶<https://plato.stanford.edu/entries/measurement-science/>

A *Sample* (UFO ECG) can be interpreted as a *Measurement* (SAREF), so we classify *Measurement* as a kind.

A crucial design decision was to classify *Time Series Measurements* as *Measurement* in SAREF, for two reasons: (i) this representation adheres to the definition of *Measurement*, i.e., measured value (*Electric Potential* units) of a property (*Heart Electrical Activity*); (ii) we reused the SAREF structure for class axioms of object properties, e.g., *hasTimestamp*, *isMeasuredIn* and *relatesToProperty*. The main implications of this choice are twofold: (1) ontologically this specialization is incorrect, since a collective cannot specialize a kind; (2) the *hasValue* property limits the value domain of a *Measurement* to exactly one float number. The *hasValues* property was added to overcome this issue, in which a *Time Series Measurements* can instantiate this property multiple times as an array of float numbers. In the conceptual model, instead of adding this specialization, the relation *makesMeasurement* between an *ECG Lead* and an *ECG Sample Sequence* makes this design decision explicit.

Finally, following the same approach of UFO ECG with *Sample Rate* data type and FHIR with *period* data property of *SampledData*, a class axiom was added in *Time Series Measurements* to relate it to a frequency (rate). For the definition of frequency, we reused SAREF4envi by importing the *hasFrequencyMeasurement* object property.

5.3. ECG device behaviour

An *ECG device* is usually referred as an ECG unit that plays the role of a *recorder* in the complex event (action) of an *ECG Recording Session*. In SAREF, we can classify this complex action as a *Task* that an ECG device *accomplishes*. Therefore, *accomplishes* plays the role of the inverse of the *hasParticipant* relationship (same as *isAccomplishedBy*). The *hasParticipant* relationship between *ECG Recording Session* and *Person Under ECG Monitoring* is represented with *dc:author* (Dublin Core).

In UFO, the *event* stereotype provides the relations *start* and *end* to limit the event temporal boundaries, as used in UFO ECG. SAREF defines the *hasTime* property, which is a “relationship to associate time information to an entity”, thus, we specialized this relation with *hasStart* and *hasEnd*, adding them as class axioms of *Task*. In UFO ECG, the other participant in this event is a *Patient*, but in SAREF4health we decided to generalize this ontological commitment, since someone does not need to be a patient (person under medical treatment) to participate in an ECG recording session. Thus, we introduced the *Person Under ECG Monitoring* element, i.e., a role of a *Living Person*, which is a phase often used in the UFO research (Guizzardi, 2005). This ontological commitment is motivated by the EWS use case, since a driver is not necessarily a patient. In SAREF4health, a *Living Person* is a *Person*, which is a term with similar definition in DOLCE Ultra-Light and schema.org. We only consider ECG for humans in this ontology.

Usually, the frequency of an *ECG device* can be set through an API, which becomes the frequency of each *ECG Sample Sequence* measured during a *Recording Session*. Therefore, we added a class axiom to the *ECG device* element: *has Frequency Measurement* property with range of only *Frequency Measurements* (from SAREF4envi, see Section 3.2). Although it seems redundant, this approach is required to differentiate the current frequency of a device from a frequency used in prior sample sequences. For example, the device Shimmer3 ECG can be set for sampling frequency (rate) of 512 Hz, which is recommended for clinical grade ECG, i.e., 512 data samples per second or an interval of 0.002 seconds between two consecutive data samples, from 0.05 Hz to 8000 Hz range. Suppose that after collecting some sample sequences and before sending the message to the gateway, the frequency is set to 256 Hz and new sample sequences are collected. With our approach, the message describes the current frequency of the device (256Hz) and the frequencies used in each collected sample sequence (512 Hz and 256 Hz).

5.4. ECG leads composition

An *ECG Device* registers the *Heart Electrical Activity* through electrodes attached to different places of the body, under the obvious assumption that the heart is beating inside the body of a living person. Two electrodes enable an *ECG lead* to be measured, which is an electrical vector characterized by the depolarization of the heart resulted by the electrical signal between the atria and the ventricles. Manufacturers commonly characterize an *ECG device* by its number of *ECG leads*. An ECG device is composed by extremity electrodes, which must be attached close to the left arm (LA), right arm (RA), left leg (LL) and the right leg (RL); and chest (precordial) electrodes, which can vary from one unit to six units (V_{1-6}). By convention, *lead I* measures the electrical activity from RA to LA, *lead II* measures of the electrical activity from RA to LL, *lead III* measures the electrical activity from LA to LL. The rule $lead I + lead III = lead II$ makes it possible to derive a lead based on the other two. *Leads I, II and III* are known as *Bipolar Limb*. *Unipolar leads* measure the electrical activity from the Wilson’s central terminal (negative pole) to each of the chest electrodes (positive poles). For example, the Shimmer3 ECG is a four-lead ECG device wired with four extremity electrodes and one chest electrode, enabling the measurement of three bipolar and one unipolar lead.

For the sake of simplicity and to avoid verbosity and follow the terminology commonly used in the industry, we decided to represent an *ECG device* according to its *ECG leads*, classifying an *ECG lead* as a *Sensor*, since a sensor “detects and responds to events or changes in the physical environment” (SAREF). A lead can be either bipolar or unipolar, and an *ECG device* consists of at least one *Bipolar Limb* and one *Unipolar* lead. An *ECG lead* measures the *Heart Electrical Activity* property and makes measurements of *ECG Sample Sequence*. Figure 4 shows the core concepts of the SAREF4health ontology discussed here in an OntoUML model.

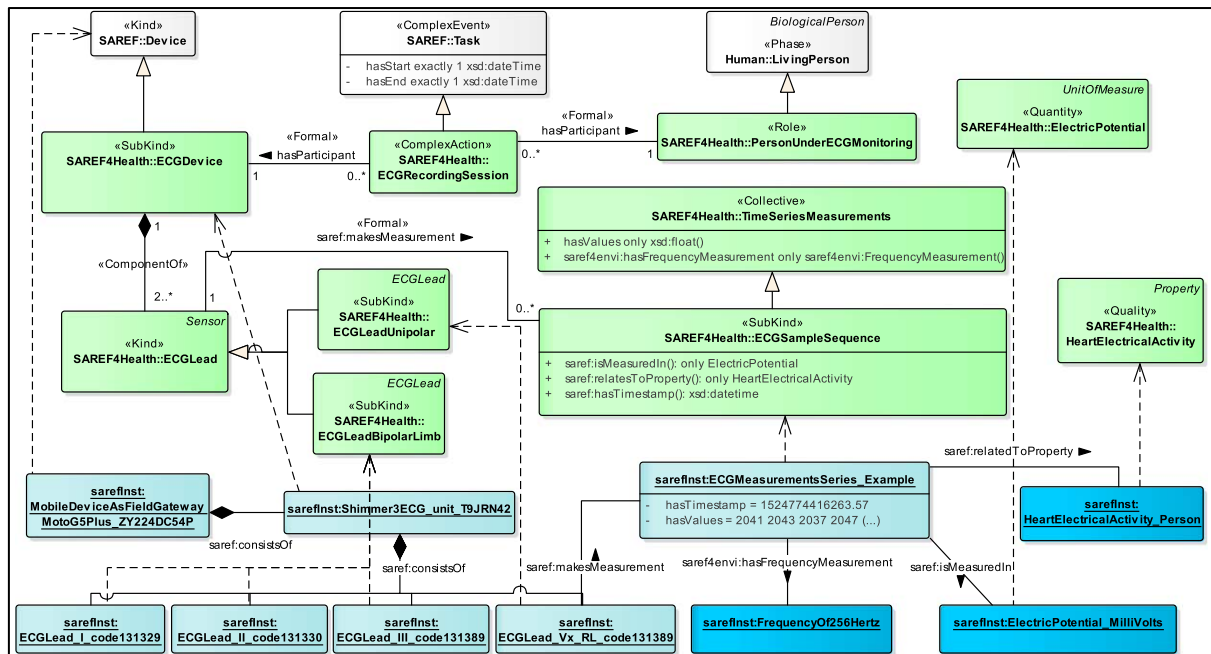


Fig. 4. SAREF4health main elements (green) with instantiated examples (blue).

6. Validation

SAREF4health was developed to address the INTER-IoT-EWS use cases, covering the requirements listed in Section 4. When developing SAREF4health, we considered the general challenges of semantic integration, semantic performance and data quality problems to provide semantic interoperability to IoT solutions. We focused on providing an ontology that balances quality, message payload for IoT scenarios and IoT orientation, as guided by the comparison study. The validation setup aimed to demonstrate how the requirements are satisfied through two activities: (1) responding to the competency questions with examples that are manually generated; and (2) the development and proper tests (functional and non-functional) of the an IoT EWS prototype, i.e., INTER-IoT-EWS²⁷ solution for semantic data streaming with standardised ontologies and JSON-LD, improving existing solutions for vehicle tracking and ECG monitoring. These validations target the evaluation of ontology coverage, i.e., whether the ontology covers the domain of interest, and ontology processing performance, i.e., whether the ontology can support the IoT use case constraints, such as total transaction time for data processing and message payload adequacy.

6.1. Responding competency questions

The first part of the SAREF4health validation was performed by specifying the ontology in RDF, using Menthor (Moreira et al., 2016) and Protégé. Instances of SAREF4health were created (as examples), and basic SPARQL queries were executed to answer each competency question listed in Table 2. As illustrated in some SAREF4health instances in Figure 4, the examples are based on the Shimmer3 ECG device, the device API (TinyOS), the mobile app (Android) as field gateway, and an IoT platform context broker playing the role of cloud gateway. Table 4 summarizes the results. Because of space limitation, we omitted textual properties (e.g., labels and comments). The SPARQL queries used are based on the following template:

```
SELECT * WHERE {{?s1 ?p1 [Name of the term]} UNION {[Name of the term] ?p2 ?o2}.}
```

All the competency questions listed in Table 2 could be answered, which is an indication for the completeness of SAREF4health and, therefore, its semantic validity. Our ontological commitments were driven by the definitions listed in Table 3. The ontology structure was assessed by analysing the SPARQL results (Table 4). For example, *time series* element (CQ05) is equivalent to *sample sequence* (UFO ECG) and *Sampled Data* (FHIR). The SPARQL results presented in Table 4 also show that the structure allows the representation of an ordered frequency-based sequence of float data, through the *hasValues only xsd:float* axiom (in CQ05), which allows the representation of multiple values and the serialisation as a compact array in JSON-LD.

SAREF4health quality was leveraged by the application of the best practices for ontology engineering adopted by our methodology. We argue that SAREF4health has a high quality because of the ontological foundations and practical matters, such as making SAREF4health as dereferenceable as possible and checking and reusing popular and common ontologies. A limitation of SAREF4health is that it still lacks a concept to represent the physician(s) administering an ECG recording session, which is common in other standards (e.g., DICOM). To overcome this issue, we suggest to add a class axiom to *ECG Recording Session* with the *dc:creator*, similarly to the approach taken with *dc:author*.

²⁷<https://github.com/jonimoreira/INTER-IoT-EWS/>

Table 4
Responding the competency questions (Table 2) with SPARQL queries

ID	Results	Description
CQ01	<pre>> rdfs:subClassOf saref:consistsOf min 1 ECGLeadUnipolar > rdfs:subClassOf saref:consistsOf min 1 ECGLeadBipolarLimb > rdfs:subClassOf saref:accomplishes only ECGRecordingSession > rdfs:subClassOf saref4envi:hasFrequencyMeasurement only saref4envi:FrequencyMeasurement > rdfs:subClassOf saref:Device < sarefInst:Shimmer3ECG_unit_T9JRN42 rdf:type</pre>	An ECG device is a device composed by at least one unipolar and one bipolar limb leads. It accomplishes the task of ECG recording session and has a specific frequency measurement.
CQ02	<pre>> rdfs:subClassOf dc:author only LivingPerson > rdfs:subClassOf saref:Task < ECGDevice saref:accomplishes only ECGRecordingSession < sarefInst:RecordingECGSession_01 rdf:type</pre>	An ECG recording session is a task in which a living person participates and is accomplished by an ECG device.
CQ03	<pre>> rdfs:subClassOf saref:measuresProperty only HeartElectricalActivity > rdfs:subClassOf saref:makesMeasurement only ECGSampleSequence > rdfs:subClassOf saref:Sensor < ECGLeadBipolarLimb rdfs:subClassOf < ECGLeadUnipolar rdfs:subClassOf</pre>	An ECG lead is a sensor that can be either bipolar limb or unipolar, which is able to measure a heart electrical activity, making measurements of the type ECG sample sequence.
CQ04	<pre>> rdfs:subClassOf saref:relatesToProperty only HeartElectricalActivity > rdfs:subClassOf saref:isMeasuredIn only ElectricPotential > rdfs:subClassOf TimeSeriesMeasurements < ECGLead saref:makesMeasurement only ECGSampleSequence < sarefInst:ECGMeasurementsSeries_Example001 rdf:type</pre>	An ECG sample sequence is a measurement time series that relates to the property of heart electrical activity, measures in an electric potential unit. ECG sample sequences are measured by ECG leads.
CQ05	<pre>> rdfs:subClassOf saref4envi:hasFrequencyMeasurement only saref4envi:FrequencyMeasurement > rdfs:subClassOf hasValues only xsd:float > rdfs:subClassOf saref:Measurement < ECGSampleSequence rdfs:subClassOf</pre>	A measurement time series is a measurement that has a frequency and a set of float values. Currently, the only type of time series measurements available is the ECG sample sequence.
CQ06	<pre>> saref4envi:FrequencyMeasurement > rdfs:subClassOf saref:Measurement > rdfs:subClassOf saref:isMeasuredIn exactly 1 saref4envi:FrequencyUnit > rdfs:subClassOf saref:relatesToProperty value saref4envi:Frequency</pre>	The frequency measurement type imported from SAREF4envi is a measurement that is measured in a frequency unit and relates to the frequency property.
CQ07	<pre>> sarefInst:Shimmer3ECG_unit_T9JRN42 rdf:type:ECGDevice; saref:consistsOf sarefInst:AccelerometerX_ECGDevice; saref:consistsOf sarefInst:AccelerometerY_ECGDevice; saref:consistsOf sarefInst:AccelerometerZ_ECGDevice; > sarefInst:AccelerometerX_ECGDevice rdf:type saref:Sensor; saref:makesMeasurement sarefInst:Measurement_AccelerationX_001; saref:measuresProperty dim:Acceleration; > sarefInst:Measurement_AccelerationX_001 rdf:type saref:Measurement; saref:hasTimestamp "2018-04-22T22:15:30"^^xsd:dateTime; saref:hasValue "100"^^xsd:float; saref:isMeasuredIn unit:metrePerSecondSquared; saref:relatesToProperty quantity:acceleration;</pre>	Tri-axial acceleration data are represented according to each accelerometer sensor (x , y , z) of the ECG device. Each accelerometer sensor measures the property acceleration and makes acceleration measurements. An acceleration measurement is measured in metre per second squared, has a value and a timestamp.
CQ08	<pre>> sarefInst:Shimmer3ECG_unit_T9JRN42 rdf:type:ECGDevice; saref4envi:hasFrequencyMeasurement sarefInst:FrequencyOf256Hertz; saref:accomplishes sarefInst:RecordingECGSession_01; saref:consistsOf sarefInst:AccelerometerX_ECGDevice; saref:consistsOf sarefInst:AccelerometerY_ECGDevice; saref:consistsOf sarefInst:AccelerometerZ_ECGDevice; saref:consistsOf sarefInst:ECGLead_III_code131389; saref:consistsOf sarefInst:ECGLead_II_code131330; saref:consistsOf sarefInst:ECGLead_I_code131329; saref:consistsOf sarefInst:ECGLead_Vx_RL_code131389; saref:consistsOf sarefInst:Shimmer3BatteryLevelSensor_T9JRN42; saref:hasManufacturer "Shimmer"; saref:hasTypicalConsumption sarefInst:Shimmer3ECGBattery;</pre>	The ECG device mereology is responsible for representing the device sensors through the <i>consists of</i> property. This structure allows (near) real-time monitoring by accumulating the measurements (made by the sensors). This includes the time series measurements (e.g., made by ECG leads) and isolated measurements (e.g., acceleration and battery level).

6.2. Semantic IoT early warning system prototype

In the second part of the validation, we connected a device (Shimmer3 ECG unit) to a field gateway (Android smartphone) to a cloud gateway (MS Azure IoT Hub). The goal of this validation is to show that SAREF4health has the adequate coverage for the envisioned INTER-IoT scenario, and also provides an adequate performance in terms of message verbosity and processing time.

The INTER-IoT-EWS prototype improves the manufacturer’s Android app (*Shimmer Xamarin Capture*) by translating the data received from the ECG device to SAREF4health and enables different rates of data exchange between device-mobile and mobile-cloud, which is a common network bandwidth optimisation requirement. This scenario allowed us to exploit the SAREF4health approach for time series measurements.

Figure 5 illustrates the solution architecture of INTER-IoT-EWS, having the MS Azure MyDriving application, which exemplifies the monitoring of vehicle trips, being able of tracking real-time location, speed and acceleration data through a mobile device connected to the cloud infrastructure (open source). The MyDriving mobile app behaves as a field gateway publishing data in the IoT Hub as JSON messages. Since the *MyDriving* app does not represent data as RDF, changes in the mobile app were required to turn the application to a *MyDriving for Linked Data* (MyDriving-LD) version. MyDriving-LD plays the role of the *Semantic Field Gateway* component of SEMIoTICS, providing logistics data, and thus requiring an ontology that represents transports, trips, location, acceleration and the goods transported.

MyDriving-LD had to be integrated with the Shimmer3 ECG device, using the *ShimmerCapture* solution as an example. For health data, we applied the SAREF4health. Finally, ontology modules produced

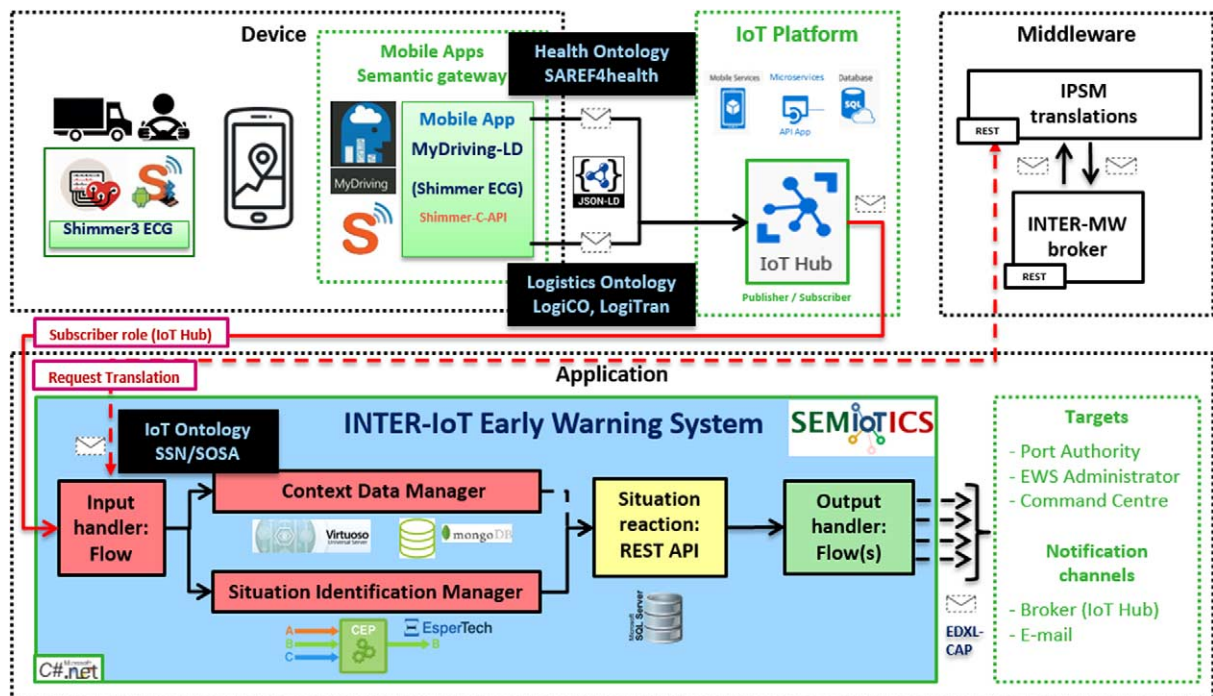


Fig. 5. An IoT EWS to detect accident risks and accidents at the port of Valencia.

by TNO were used for logistics data²⁸: (1) *Logistics Core Ontology* (LogiCO) ontology, along with its extensions *Transport* (LogiTrans) and *Logistics Services* (LogiServ) ontologies.

The Shimmer ECG 3 device collects ECG data from the truck driver and sends the data to the MyDriving-LD application via *Bluetooth* at a high-frequency in real time. MyDriving-LD is deployed in a smartphone and plays the role of a field gateway, also managing the logistics data. A semantic wrapper is implemented in MyDriving-LD to allow the annotation of aggregated data with SAREF4health and the logistics ontologies. These annotated data are sent to the cloud as JSON-LD messages, i.e., published in the cloud gateway (*Azure IoT Hub*).

In the main configuration, the INTER-IoT-EWS *Input Handler* subscribes to receive all messages published in *Azure IoT Hub*, checking each message type (the domain: IoT, health or logistics) and certifying whether translations to harmonize the data in the SEMIoTICS core ontology are necessary. If this is true, the *Input Handler* requests the translations to the INTER-IoT semantic mediator (IPSM) (Ganzha et al., 2017a), forwarding the input message and the alignment required (pre-configured) for executing the translations. The IPSM is responsible for syntactically and semantically translating the data streams. After the semantic translations execution, the *Input Handler* is responsible for the syntactical translation to internal POCO classes that follow the SEMIoTICS core context model, which are used for both *Context Data* and *Situation Identification* managers. Once a situation is identified, the *Situation Reaction* component (see Figure 5) checks the workflows (business processes) to be triggered, which describe the different targets and their information requirements. The output handler formats the data according to the SEMIoTICS notification ontology (from OASIS EDXL-CAP) and forwards the messages with the appropriate information for each target.

The validation plan checked whether the requirements (see Section 4) were fulfilled. The functional validation activities included factory acceptance tests (FAT) and site acceptance tests (SAT) through a pilot in the port of Valencia, where accidents were simulated in accordance with the port emergency exercises. Both FAT and SAT assess whether the system works for the intended risks' detection and warnings. The FAT was performed to check completeness and requirements verification, simulating each of the situation types of each use case. Test hooks included data injection to increase the instantaneous acceleration data for the collision detection and to change the processed heart rate to below the bradycardia threshold and above the tachycardia threshold, as well as other more complex data patterns. An UI prototype was developed to support the execution of the tests.

The overall INTER-IoT-EWS performance analysis included the measurement of the total transaction time for the whole EWS workflow. We followed a similar validation method of the *Semantic IoT EWS* approach for scalability and resilience, testing upstream data acquisition and situation identification, measuring total transaction time from one component to another. The validation also considered the multi-brokering multi-cluster approach, showing that an industry-oriented cloud infrastructure, such as MS Azure, is able to scale the throughput threshold of 700 msg/sec/unit from the TRIDEC project (Middleton et al., 2013). In terms of efficiency of publishing semantically rich sensor data with a multi-broker approach, we conclude that this threshold (700 msg/sec/unit), considering a message size of 6 Kb (4.200 Kb/sec/unit), is easily reached with the Azure cloud infrastructure. The IoT Hub quotas and throttling supports the throughput of 6.000 msg/sec/unit ("device-to-cloud"), considering a message size of 4 Kb (24.000 Kb/sec/unit). Therefore, the cloud infrastructure can scale up to five times the broker throughput of the original approach (TRIDEC).

²⁸<http://ontology.tno.nl/logico/>, <http://ontology.tno.nl/transport/>, <http://ontology.tno.nl/logiserv/>

7. Discussion

An important contribution of this paper is the proposed methodology to deal with both conceptual modelling and implementation concerns. We argue that SAREF4health has high ontological expressiveness because of its completeness for the ECG domain, since “the notion of completeness at the level of individual specifications is related to the notion of ontological expressiveness” (Guizzardi, 2005). Moreover, we claim that construct redundancy (Bera and Geert, 2017) was eliminated from SAREF4health regarding the elements related to *Observation* from UFO ECG. Although these terms are ontologically correct, in UFO ECG they bring additional complexity, which is represented implicitly in SAREF through the *make measurement* property and the *Measurement* class. While in SAREF a device makes a measurement, which has a timestamp, a value and is measured in a specific unit; in UFO ECG a device carries an observation, which has a timestamp, and this observation produces a sample, which has a value and is measured in a specific unit.

As an additional contribution of this work, clear mappings between UFO ECG to FHIR were identified, as (UFO) *Observation series* to (FHIR) *Observation*, (UFO) *Sample sequence* to (FHIR) *Sampled Data*, (UFO) *Sample sequence > sample-sequence-of* is the inverse of (FHIR) *Observation > valueSampledData*. Other mappings can be easily extracted from the results of our ontological analysis described in Section 4.2. Regarding the use of *hasValues* with *TimeSeriesMeasurements* for ordered values, this approach works well when serializing with the *@list* element of JSON-LD, but poses a formalization issue. Although the RDF language syntax provides the *rdf:Seq* element for ordered lists, ordering triples is not covered by the RDF model, which impacts on the use of OWL2-DL reasoners. The SAREF4health ontology (TTL) and the prototype are available for download²⁹.

Although SAREF4Health and FHIR RDF seem competitors, they are complementary. They overlap in their goals to some extent, i.e., both target interoperability improvement of healthcare data exchange. While FHIR defines resource templates for clinical information modelling, e.g. *Patient*, *Practitioner*, *Organization*, *Healthcare Service*, *Clinical*, *Diagnostics*, *Medications*, *Workflow* and *Financial* terms, SAREF is oriented to measurements, measurement processes/sessions and medical devices’ modelling, e.g., sensors’ assembly. FHIR is developed for providing reasoning and decision making support for healthcare processes and clinical systems, more at the organizational level, while SAREF4Health is at the engineering level. Independent of these differences, we believe that this work provides a relevant contribution towards a FHIR ontology rather than a straightforward RDF representation of the FHIR lexicon data model. We recommend that, at least, FHIR RDF object properties are harmonized through an ontological analysis of the data model. Although the FHIR lexicon standard is widely adopted by industry (e.g., Google and Apple), its RDF version is not yet mature.

Although overlapping, the scope of STF 556 (SAREF extensions for SAREF4ehaw and SAREF4wear) is broader than the validation scope of the SAREF4health (ECG monitoring). The STF is targeting general requirements from different stakeholders in other use cases. The STF is adopting the more general terms from SAREF4health, e.g., time series measurements, recording session and person under monitoring. SAREF4health is being studied for SAREF4wear for the composition of leads in ECG devices and electrodes to be attached to a person’s chest.

²⁹<https://github.com/jonimoreira/INTER-IoT-EWS>

8. Conclusions

Current IoT-based EWS approaches partially address the semantic integration of a variety of data sources along with processing in time-critical applications and data analysis for effective responses. Existing standardised ontologies for IoT are ETSI SAREF and W3C SSN/SOSA, which provide the benefit of a reference interoperability language for IoT, but are quite verbose when representing time series data in common healthcare scenarios, as for ECG monitoring. In this paper, we proposed SAREF4health, which extends SAREF and combines ontology-driven conceptual modelling, standardisation practices and RDF implementation of stream data. SAREF4health was supported by the SEMIoTICS framework and is able to represent real-time ECG time series of sensor measurements that are exchanged between the field (mobile device) and the cloud (context broker) gateways.

Our validation showed that a trade-off between ontology quality and lightweight data serialisation was a crucial aspect of SAREF4health design. The use of an ECG reference ontology grounded in UFO theory played a major role to improve semantic quality. Furthermore, the reuse of standardised ontologies showed to be essential to better understand both the IoT and healthcare domains, as well as implementation constraints. SAREF4health addressed the performance requirement of time series measurements in the message size (payload), i.e., the verbosity problem, validated through a prototype of an ontology-driven health IoT platform, which transmits data from the device to a mobile to the cloud in different frequencies.

Future work includes formal empirical validation to compare SAREF4health with other related ontologies, especially HL7 FHIR RDF. In addition, semantic translations between SAREF4health and FHIR RDF should be developed to address the integration requirements for different IoT platforms that support data acquisition. Semantic loss and processing impact of these translations should be measured to evaluate the mappings in order to identify bottlenecks.

References

- Bera, P. & Geert, P. (2017). The effects of construct redundancy on readers' understanding of conceptual models. *Journal of Database Management*. doi:[10.4018/JDM.2017070101](https://doi.org/10.4018/JDM.2017070101).
- Bond, R.R., Finlay, D.D., Nugent, C.D. & Moore, G. (2011). A review of ECG storage formats. *International Journal of Medical Informatics*, 80(10), 681–697. doi:[10.1016/j.ijmedinf.2011.06.008](https://doi.org/10.1016/j.ijmedinf.2011.06.008).
- Brandt, P., Basten, T., Stuiik, S., Bui, V., Clercq, P.d., Pires, L.F. & Sinderen, M.v. (2013). Semantic interoperability in sensor applications making sense of sensor data. In *2013 IEEE Symposium on Computational Intelligence in Healthcare and e-Health (CICARE)*, 16–19 April 2013. doi:[10.1109/CICARE.2013.6583065](https://doi.org/10.1109/CICARE.2013.6583065).
- Cosío-León, M.A., Ojeda-Carreño, D., Nieto-Hipólito, J.I. & Ibarra-Hernández, J.A. (2018). The use of standards in embedded devices to achieve end to end semantic interoperability on health systems. *Computer Standards & Interfaces*, 57, 68–73. doi:[10.1016/j.csi.2017.11.006](https://doi.org/10.1016/j.csi.2017.11.006).
- da Costa, C.A., Pasluosta, C.F., Eskofier, B., da Silva, D.B. & da Rosa Righi, R. (2018). Internet of health things: Toward intelligent vital signs monitoring in hospital wards. *Artificial Intelligence in Medicine*. doi:[10.1016/j.artmed.2018.05.005](https://doi.org/10.1016/j.artmed.2018.05.005).
- Daniele, L., den Hartog, F. & Roes, J. (2015). Created in close interaction with the industry: The smart appliances REFERENCE (SAREF) ontology. In *Formal Ontologies Meet Industry: 7th International Workshop (FOMI)*. doi:[10.1007/978-3-319-21545-7_9](https://doi.org/10.1007/978-3-319-21545-7_9).
- Daniele, L., den Hartog, F. & Roes, J. (2016). How to keep a reference ontology relevant to the industry: A case study from the smart home. In *Ontology Engineering*, Cham. doi:[10.1007/978-3-319-33245-1_12](https://doi.org/10.1007/978-3-319-33245-1_12).
- EC (2017). Rolling plan for ICT standardisation 2017. Retrieved on April 2, 2020 from <https://ec.europa.eu/digital-single-market/en/news/rolling-plan-ict-standardisation>.
- ETSI (2019a). Smart Body Area Network (SmartBAN): Unified data representation formats, semantic and open data model. ETSI Technical Specification (TS) 103 378 V1.1.1 (2015-12).
- ETSI (2019b). Smart Body Area Networks (SmartBAN): Service and application standardized enablers and interfaces, APIs and infrastructure for interoperability management. ETSI Technical Report (TR) 103 394 V1.1.1 (2018-01).

- Fernandez, J.d.R., Toma, D.M., Martinez, E., Jirka, S. & O'Reilly, T. (2019). Chapter 6 – from sensor to user: Interoperability of sensors and data systems. In: *Challenges and Innovations in Ocean in Situ Sensors* (pp. 289–337). Elsevier. doi:[10.1016/B978-0-12-809886-8.00006-5](https://doi.org/10.1016/B978-0-12-809886-8.00006-5).
- Fortino, G. & Gravina, R. (2015). Fall-MobileGuard: A smart real-time fall detection system. In *Proceedings of the 10th EAI International Conference on Body Area Networks*, Sydney, New South Wales, Australia. doi:[10.4108/eai.28-9-2015.2261462](https://doi.org/10.4108/eai.28-9-2015.2261462).
- Fortino, G., Parisi, D., Pirrone, V. & Di Fatta, G. (2014). BodyCloud: A SaaS approach for community body sensor networks. *Future Generation Computer Systems*, 35, 62–79. doi:[10.1016/j.future.2013.12.015](https://doi.org/10.1016/j.future.2013.12.015).
- Ganzha, M., Paprzycki, M., Pawłowski, W., Szmeja, P. & Wasielewska, K. (2016). Semantic technologies for the IoT – an inter-IoT perspective. In *2016 IEEE First International Conference on Internet-of-Things Design and Implementation (IoTDI)*, 4–8 April 2016. doi:[10.1109/IoTDI.2015.22](https://doi.org/10.1109/IoTDI.2015.22).
- Ganzha, M., Paprzycki, M., Pawłowski, W., Szmeja, P. & Wasielewska, K. (2017a). Streaming semantic translations. In *2017 21st International Conference on System Theory, Control and Computing (ICSTCC)*, 19–21 Oct. 2017. doi:[10.1109/ICSTCC.2017.8107003](https://doi.org/10.1109/ICSTCC.2017.8107003).
- Ganzha, M., Paprzycki, M., Pawłowski, W., Szmeja, P., Wasielewska, K. & Palau, C.E. (2017b). From implicit semantics towards ontologies — practical considerations from the INTER-IoT perspective. In *2017 14th IEEE Annual Consumer Communications & Networking Conference (CCNC)*, 8–11 Jan. 2017. doi:[10.1109/CCNC.2017.7983082](https://doi.org/10.1109/CCNC.2017.7983082).
- Gonçalves, B., Guizzardi, G. & Pereira Filho, J.G. (2011). Using an ECG reference ontology for semantic interoperability of ECG data. *Journal of Biomedical Informatics*, 44(1), 126–136. doi:[10.1016/j.jbi.2010.08.007](https://doi.org/10.1016/j.jbi.2010.08.007).
- Gonçalves, B., Zamborlini, V. & Guizzardi, G. (2009). An ontological analysis of the electrocardiogram. *Electronic Journal of Communication, Information and Innovation in Health*, 1–26. doi:[10.3395/reciis.v3i1.242en](https://doi.org/10.3395/reciis.v3i1.242en).
- Guizzardi, G. (2005). Ontological foundations for structural conceptual models. Enschede. Retrieved on April 2, 2020 from <http://doc.utwente.nl/50826/>.
- Gyrard, A., Serrano, M. & Atemezing, G.A. (2015). Semantic web methodologies, best practices and ontology engineering applied to Internet of Things. In *2015 IEEE 2nd World Forum on Internet of Things (WF-IoT)*, 14–16 Dec. 2015. <https://doi.ieeecomputersociety.org/10.1109/WF-IoT.2015.7389090>.
- Haghi, M., Thurow, K. & Stoll, R. (2017). Wearable devices in medical Internet of things: Scientific research and commercially available devices. *Healthcare Informatics Research*, 23(1), 4–15. doi:[10.4258/hir.2017.23.1.4](https://doi.org/10.4258/hir.2017.23.1.4).
- Heiler, S. (1995). Semantic interoperability. *ACM Comput. Surv.*, 27(2), 271–273. doi:[10.1145/210376.210392](https://doi.org/10.1145/210376.210392).
- HL7 (2005). HL7 Version 3 Implementation Guide: Annotated ECG (aECG) R1, Release 2. Retrieved on April 2, 2020 from https://www.hl7.org/implement/standards/product_brief.cfm?product_id=102.
- Hossain, M.S. & Muhammad, G. (2016). Cloud-assisted industrial Internet of Things (IIoT) – enabled framework for health monitoring. *Computer Networks*, 101, 192–202. doi:[10.1016/j.comnet.2016.01.009](https://doi.org/10.1016/j.comnet.2016.01.009).
- Khodadadi, F. & Sinnott, R.O. (2017). A semantic-aware framework for service definition and discovery in the Internet of Things using CoAP. *Procedia Computer Science*, 113, 146–153. doi:[10.1016/j.procs.2017.08.334](https://doi.org/10.1016/j.procs.2017.08.334).
- Kim, J., Kang, S., Lee, J. & Choi, B.W. (2012). A semantic translation method for data communication protocols. *Journal of Systems and Software*, 85(12), 2876–2898. doi:[10.1016/j.jss.2012.06.018](https://doi.org/10.1016/j.jss.2012.06.018).
- Li, H., Seed, D., Flynn, B., Mladin, C. & Girolamo, R.D. (2016). Enabling semantics in an M2M/IoT service delivery platform. In *Proceedings – 2016 IEEE 10th International Conference on Semantic Computing (ICSC)* (pp. 206–213). doi:[10.1109/ICSC.2016.28](https://doi.org/10.1109/ICSC.2016.28).
- Middleton, S., Zlatev, Z., Mossgraber, J., Chaves, F. & Tao, R. (2013). The seven main challenges of an early warning system architecture. In *Information Systems for Crisis Response and Management (ISCRAM)*. Retrieved on April 2, 2020 from <https://eprints.soton.ac.uk/359361/>.
- Moreira, J.L.R., Daniele, L.M., Pires, L.F., Sinderen, M.V., Wasielewska, K., Szmeja, P. et al. (2017a). Towards IoT platforms' integration: Semantic translations between W3C SSN and ETSI SAREF. In *Semantics. Workshop SIS-IoT*. Retrieved on April 2, 2020 from <http://ceur-ws.org/Vol-2063/sisiot-paper3.pdf>.
- Moreira, J.L.R., Ferreira Pires, L., Sinderen, M.V. & Costa, P.D. (2017b). Ontology-driven conceptual modeling for early warning systems: Redesigning the situation modeling language. In *MODELSWARD*. doi:[10.5220/0006208904670477](https://doi.org/10.5220/0006208904670477).
- Moreira, J.L.R., Ferreira Pires, L., Sinderen, M.v. & Dockhorn Costa, P. (2015a). Developing situation-aware applications for disaster management with a distributed rule-based platform. In *RuleML 2015 Doctoral Consortium*, Berlin, Germany. Retrieved on April 2, 2020 from <http://doc.utwente.nl/97411/>.
- Moreira, J.L.R., Ferreira Pires, L., Sinderen, M.v. & Dockhorn Costa, P. (2015b). Towards ontology-driven situation-aware disaster management. *Journal of Applied Ontology*, 10(3–4), 339–353. doi:[10.3233/AO-150155](https://doi.org/10.3233/AO-150155).
- Moreira, J.L.R., Pires, L.F. & Sinderen, M.V. (2018a). SAREF4health: IoT standard-based ontology-driven healthcare systems. In *10th International Conference on Formal Ontology in Information Systems, (FOIS 2018)*, Cape Town, South Africa. Best paper award. doi:[10.3233/978-1-61499-910-2-239](https://doi.org/10.3233/978-1-61499-910-2-239).
- Moreira, J.L.R., Pires, L.F. & Sinderen, M.V. (2018b). Semantic interoperability for the IoT: Analysis of JSON for linked data. In *Enterprise Interoperability IX: I-ESA Proceedings. BD4EI Workshop*. doi:[10.1002/9781119564034.ch20](https://doi.org/10.1002/9781119564034.ch20).

- Moreira, J.L.R., Pires, L.F., Sinderen, M.V., Wieringa, R., Singh, P., Costa, P.D. & Llop, M. (2018c). Improving the semantic interoperability of IoT early warning systems: The port of Valencia use case. In *Enterprise Interoperability IX: I-ESA Proceedings*. doi:[10.1007/978-3-030-13693-2_2](https://doi.org/10.1007/978-3-030-13693-2_2).
- Moreira, J.L.R., Sales, T.P., Guerson, J., Braga, B.F.B., Brasileiro, F. & Sobral, V. (2016). Mentor editor: An ontology-driven conceptual modeling platform. In *9th International Conference on Formal Ontology in Information Systems (FOIS 2016) – Demonstration Track*, Annecy, France. Retrieved on April 2, 2020 from <http://ceur-ws.org/Vol-1660/demo-paper1.pdf>.
- Nachabe Ismail, L., Girod-Genet, M. & El Hassan, B. (2016). Semantic techniques for IoT data and service management: ONTOSMART system. *International Journal of Wireless & Mobile Networks (IJWMN)*, 8(4), 43–63. doi:[10.5121/ijwmn.2016.8403](https://doi.org/10.5121/ijwmn.2016.8403).
- Olivieri, A.C., Bocchi, Y. & Rizzo, G. (2016). Chapter 4 – integration in the Internet of Things: A semantic middleware approach to seamless integration of heterogeneous technologies. In C. Dobre and F. Xhafa (Eds.), *Pervasive Computing* (pp. 97–128). Boston: Academic Press. doi:[10.1016/B978-0-12-803663-1.00004-8](https://doi.org/10.1016/B978-0-12-803663-1.00004-8).
- Palavalli, A., Karri, D. & Pasupuleti, S. (2016). Semantic Internet of Things. In *IEEE Tenth International Conference on Semantic Computing (ICSC)*.
- Rahmani, A.M., Gia, T.N., Negash, B., Anzanpour, A., Azimi, I., Jiang, M. & Liljeberg, P. (2018). Exploiting smart e-health gateways at the edge of healthcare Internet-of-Things: A fog computing approach. *Future Generation Computer Systems*, 78, 641–658. doi:[10.1016/j.future.2017.02.014](https://doi.org/10.1016/j.future.2017.02.014).
- Rezaei, R., Chiew, T.K. & Lee, S.P. (2014). An interoperability model for ultra large scale systems. *Advances in Engineering Software*.
- Scherp, A., Saathoff, C., Franz, T. & Staab, S. (2011). Designing core ontologies. *Journal of Applied Ontology*, 6, 177–221. doi:[10.3233/AO-2011-0096](https://doi.org/10.3233/AO-2011-0096).
- Solbrig, H.R., Prud, E., Grieve, G., McKenzie, L., Mandel, J.C., Sharma, D.K. & Jiang, G. (2017). Modeling and validating HL7 FHIR profiles using semantic web shape expressions (ShEx). *Journal of Biomedical Informatics*, 67, 90–100. doi:[10.1016/j.jbi.2017.02.009](https://doi.org/10.1016/j.jbi.2017.02.009).
- Szilagyi, I. & Wira, P. (2016). Ontologies and semantic web for the Internet of Things – a survey. In *IECON 2016 – 42nd Annual Conference of the IEEE Industrial Electronics Society*, 23–26 Oct. 2016.
- Xia, H., Asif, I. & Zhao, X. (2013). Cloud-ECG for real time ECG monitoring and analysis. *Computer Methods and Programs in Biomedicine*, 110(3), 253–259. doi:[10.1016/j.cmpb.2012.11.008](https://doi.org/10.1016/j.cmpb.2012.11.008).