A Semantic Architecture for Continuous Health Monitoring, Risk Prediction, and Proactive Decision Making

Mbithe Nzomo[0000−0002−2923−8333] and
Deshendran Moodley[0000−0002−4340−9178]

Abstract Wearable sensors combined with health records and user-submitted data are becoming ubiquitous for continuous health monitoring. Semantic web technologies are well suited for representing and reasoning over these heterogeneous health data. However, existing semantic health monitoring architectures have notable deficiencies, particularly in interoperability, situation prediction, and uncertainty handling. We propose a semantic architecture that integrates an ontology with rules, fuzzy inference, and machine learning to detect and predict health risks using heterogeneous health data. We illustrate its application through a use case of atrial fibrillation and demonstrate its ability to detect and predict health situations, as well as provide decision support aligned with established health workflows and clinical guidelines.

Key words: Semantic architecture, Ontologies, Rule-based reasoning, Fuzzy inference, Health monitoring

1 Introduction

The increasing prevalence of non-communicable diseases has given rise to the emerging field of precision health, which focuses on assessing individual circumstances for early detection, prevention, and mitigation of diseases. A key aspect of this is

M. Nzomo (✉)
Department of Computer Science, University of Cape Town, South Africa
Centre for Artificial Intelligence Research (CAIR), South Africa
e-mail: mnzomo@cs.uct.ac.za

D. Moodley
Paris Institute for Advanced Study, Paris, France (2023-2024 Fellow)
Department of Computer Science, University of Cape Town, South Africa
Centre for Artificial Intelligence Research (CAIR), South Africa
e-mail: deshen@cs.uct.ac.za

incorporating continuous health monitoring into people's daily lives outside clinical settings using wearable sensors [11]. However, sensor observations alone are insufficient for health monitoring since non-communicable diseases are influenced by demographics, medical history, and lifestyle factors. This additional health data can be derived from health records and questionnaires, and can be interpreted using expert medical knowledge.

Semantic web technologies are well-established for the representation of heterogeneous data and provide powerful reasoning and deliberation capabilities. The most prominent of these are ontologies (knowledge bases of domain concepts and their relationships specified in a knowledge representation language based on formal logic [15]) and knowledge graphs (knowledge bases structured in a graph [34]). In previous work, we undertook a mapping study of the use of semantic web technologies in sensor-based personal health monitoring systems [24]. We identified key challenges that semantic architectures must address: interoperability, situation analysis (the detection and prediction of health situations), decision support, context awareness, and uncertainty handling. We critically evaluated the state of the art in the field and determined the extent to which these challenges have been addressed in current semantic systems. The study showed notable deficiencies in existing work, particularly in the areas of interoperability, situation analysis, and uncertainty handling.

A key aspect of interoperability that is poorly explored in existing semantic architectures is process and clinical interoperability, which entails a shared understanding of healthcare processes and the seamless transfer of care between different clinical teams [2]. This involves supporting clinicians in providing consistent care outside clinical settings that is grounded in established clinical guidelines and workflows. With regard to situation analysis, several semantic health monitoring systems are capable of detecting abnormalities in sensor observations or classifying individuals into predefined health states. However, many are unable to predict the risk of adverse outcomes in the future. Additionally, many do not take uncertainty into consideration in the situation analysis or decision support processes. While these challenges have been tackled outside semantic systems, there is a need for semantic architectures that integrate techniques for effective situation prediction and uncertainty handling, while supporting decision making for health and wellness applications.

To address the shortcomings of current semantic personal health monitoring systems, we propose a semantic architecture for health monitoring using heterogeneous health data. We demonstrate and evaluate the architecture through a use case application of atrial fibrillation (AF), the most common sustained arrhythmia. We use a top-down approach to create a knowledge graph by first developing an AF monitoring application ontology to semantically and formally represent health data, and then using the ontology to perform inferences on the data [34]. Although ontologies excel in structuring domain knowledge, they lack explicit support for uncertainty [23]. To deal with this, we incorporate rules supported by fuzzy inference to reason over the data for situation analysis and decision support under uncertainty. Rule-based reasoning is intuitive and human-readable, resulting in high transparency and explainability [13], while fuzzy inference is useful for expressing concepts in vague terms, which addresses the inherent uncertainty in health situations and outcomes.

We also show how outputs from machine learning (ML) classification models can be incorporated in the knowledge graph to enhance situation analysis. As illustrated by the use case, the proposed architecture is capable of detecting and predicting adverse health situations, and offering appropriate decision support firmly rooted in well-established clinical guidelines and workflows.

2 Related Work in Sensor-Based Personal Health Monitoring Systems

Semantic systems The use of semantic web technologies is ubiquitous in the health domain. However, a notable limitation of these technologies is their inability to inherently reason under uncertainty. It is therefore essential to augment semantic architectures with techniques that support uncertainty handling. One such technique is Bayesian networks (BNs), probabilistic graphical models in the form of directed acyclic graphs that represent cause-effect relationships. BNs have been used in semantic health monitoring systems to represent causes and risk factors of diseases, as is done by Kordestani et al. [19] and Mcheick et al. [22]. While BNs are suitable for modelling uncertainty in causality, they are not conducive to modelling uncertainty in raw or derived health data.

This challenge can be overcome using fuzzy inference, the process of mapping crisp input into imprecise output. Fuzzy inference is capable of converting crisp health data values into degrees of membership of defined health categories. This allows health data to be expressed in fuzzy boundaries, which is representative of human reasoning. This has been incorporated in a number of semantic health monitoring systems. For instance, Ali et al. [1] fuzzify sensor data such as blood pressure and heart rate, while Esposito et al. [8] fuzzify the intensity of physical activity. Despite several semantic personal health monitoring systems incorporating these uncertainty handling techniques, it is often only done to a limited extent. Most of these systems do not account for uncertainty across the end-to-end health monitoring process from data pre-processing to situation analysis and decision support.

Situation prediction is similarly limited in current semantic personal health monitoring systems. While a number of existing semantic systems anticipate adverse health outcomes (for example, Hristoskova et al. [17] determine the risk of congestive heart failure over a four-year time horizon, while De Brouwer et al. [6] anticipate headache attacks based on triggers), most focus on the detection of current health conditions without considering future outcomes. Our mapping study showed that only 13 out of 40 selected systems achieve some degree of situation prediction [24]. Additionally, many of these systems demonstrate ad hoc analyses that lack a solid foundation in established healthcare protocols. Thus, many existing semantic architectures fall short on building up on standardized healthcare workflows and guidelines. It is essential to develop semantic systems that facilitate collaboration between monitored individuals and clinicians, emphasizing established healthcare processes with humans in the loop.

Non-semantic systems There also exist many sensor-based personal health monitoring systems that do not incorporate semantic web technologies. A significant number of these non-semantic systems incorporate predictive models for outcomes ranging from cardiovascular disease [18], mortality and readmissions [5], and clinical laboratory measurements [7]. Among these systems, data-driven approaches, particularly deep learning, remain dominant. These models are well suited to learning complex features in dynamic sensor data. However, interpretability, explainability, computational expense, and the need for large amounts of data remain challenges in training and deploying neural networks [16]. In contrast, semantic web technologies are known to be more human-understandable and do not rely on large amounts of training data. There is a need to develop semantic architectures that incorporate ML, where the strengths of data-driven and knowledge-driven approaches can be leveraged to achieve the best of both worlds.

AF monitoring systems The most widely accepted signal for the diagnosis of AF is the electrocardiogram (ECG), a record of the heart's electrical activity. ECG data has been used extensively in AF monitoring systems for both situation detection and prediction [32]. While ECG remains the gold standard, some alternative physiological signals to detect AF have been proposed. The most prevalent of these is photoplethysmography (PPG), an optical sensing technology consisting of a light-emitting diode and a photodetector to detect blood volume changes. PPG is low-cost and non-intrusive, and has been used in a number of systems with promising results [25]. However, the accuracy of the signal can be affected by sensor placement, motion artifacts, contact pressure, skin tone, and obesity [28]. This can be alleviated by combining ECG and PPG, which has shown significant potential for AF monitoring [30].

Additionally, demographic (e.g. age and sex) and anthropometric (e.g. height and weight) data, medical history, and lifestyle factors are all important aspects of AF monitoring due to their impact on the likelihood of developing AF and its associated health risks. This data can be queried from health record databases. For example, Feldman et al. [9] integrate Apple Watch data with health records to determine AF patient eligibility for anticoagulation therapy given the risk of stroke. Non-sensor data can also be solicited directly from individuals via questionnaires. This can be particularly useful for describing AF symptoms and their severity based on impact on quality of life. For instance, the Atrial Fibrillation Effect on QualiTy-of-life (AFEQT) questionnaire was developed to assess the impact of AF [33].

3 Semantic Architecture

Design and Approach The design of the proposed architecture is guided by the key challenges identified in our previous mapping study. An important commonality among these challenges is the need for representation support. To define these challenges as design goals, we adapt the concept of competency questions (CQs), which are natural language questions typically used for requirements definition in ontology engineering [29]. We formulate ten design goals expressed as CQs, which are then used to analyse the architecture. The CQs are shown in Table 1. Note that the focus of this work is on a prototype semantic architecture for continuous monitoring, situation prediction, and decision making. Non-functional requirements such as privacy and security, while important issues in a real-world monitoring system, are out of scope of this study. A running use case application of AF is used to iteratively design and evaluate the architecture. This is described in detail in Section 4.

Table 1: The requirements of the architecture expressed as competency questions.

Challenge	CQ#	Competency Question
Interoperability	CQ1	Does the architecture support process/clinical interoperability (i.e. a common understanding of healthcare processes and the seamless transfer of care between different clinical teams)?
	CQ2	Does the architecture support semantic interoperability (i.e. the definition of health data in an unambiguous and universally understandable way)?
	CQ3	Does the architecture support syntactic interoperability (i.e. the representation of health data in a standard structure and syntax)?
Context awareness	CQ4	Does the architecture support the representation of contextual information such as identity, location, time, and activity?
Situation analysis	CQ5	Does the architecture support techniques for both the detection of current health situations and the prediction of future health situations?
Decision support	CQ6	Does the architecture support techniques that enhance transparency and explainability?
	CQ7	Does the architecture support the generation of recommendations and warnings in response to detected and predicted situations?
Uncertainty handling	CQ8	Does the architecture support data pre-processing techniques to handle missing, noisy, or otherwise invalid data?
	CQ9	Does the architecture support uncertainty handling in both the situation analysis and decision support processes?
Cross-cutting	CQ10	Does the architecture support the representation of sensor observations, personal information, and expert medical knowledge?

Abstract Architecture To fulfill the design goals, we propose an abstract semantic architecture consisting of three layers as shown in Fig. 1.

i. *Data Layer:* This layer supports the pre-processing of heterogeneous health data, i.e. sensor observations, health records, and data from questionnaires. Importantly, all data is formatted by mapping it to the Health Level 7 (HL7) Fast Health Interoperability Resources (FHIR)[1] standard for health data exchange.
ii. *Analysis Layer:* This layer handles the situation analysis functionality of the architecture using ML, rules, and fuzzy inference. This entails analysing health

[1] https://build.fhir.org/

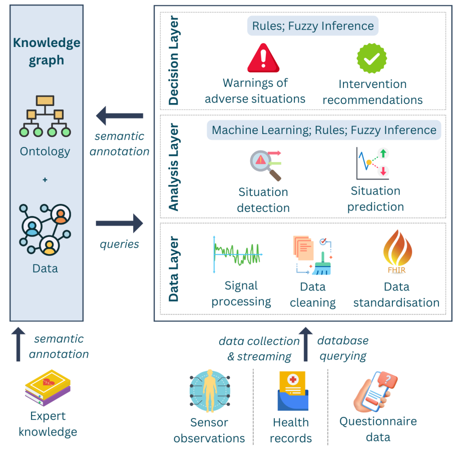

Fig. 1 Abstract semantic architecture.

data to detect current health situations, while also predicting future health situations based on current data.

iii. *Decision Layer:* Central to this layer is the integration of well-established healthcare workflows which are dynamically updated to align with the latest clinical guidelines. Adhering to these guidelines, the layer facilitates decision making in response to both detected and predicted health situations by generating recommendations, issuing warnings for adverse conditions, and providing pertinent alerts for both clinicians and monitored individuals.

A cross-cutting component of the three layers is the knowledge graph, in which all data, knowledge, situations, recommendations, and alerts are semantically annotated, represented, and stored.

4 Use Case Application: Atrial Fibrillation

AF presents a significant global health burden, affecting an estimated 46.3 million individuals worldwide as of 2016 and increasing steadily in incidence and prevalence [20]. AF symptoms include chest pain, heart palpitation, shortness of breath, and fatigue. However, it is sometimes asymptomatic, highlighting the need for early detection through continuous monitoring.

The AF monitoring process detailed in this section is anchored in the most recent European [14] and Australian [4] guidelines for AF diagnosis and management. Fig. 2 shows this process in a flow diagram. Facilitating collaboration between monitored individuals and clinicians, the process outlines clinician-led, individual-led, and shared interventions in response to detected and predicted situations. Thus, the architecture complements established AF workflows, ensuring active human participation throughout the monitoring process. In the rest of this section, we analyse the AF monitoring process following the three layers of the architecture.

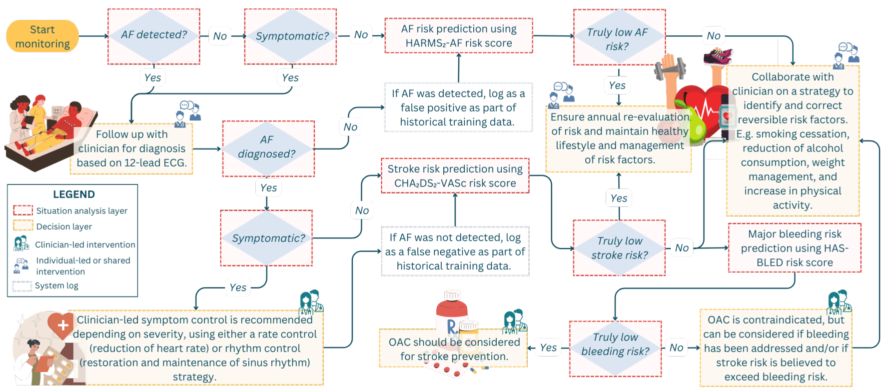

Fig. 2: Flow diagram showing the AF monitoring process. ECG: electrocardiogram. OAC: oral anticoagulation.

I. Data Layer In this layer, health data relevant to AF is preprocessed. Although live sensor data was not used in this study, a wearable device such as the Polar H10 chest strap[2] can be used to collect ECG data. This device is equipped with Bluetooth Low Energy technology, allowing for extraction of the raw single-lead ECG data using a library such as Python's Bluetooth Low Energy platform Agnostic Klient (Bleak)[3]. The ECG data can be encoded as a FHIR Observation resource. Techniques such as filtering, detection of waveform characteristics, frequency conversion, and normalization can be used to preprocess the data in preparation for situation analysis.

The architecture assumes that comprehensive health records, including demographic and anthropometric data as well as diagnosis, procedure, and medication history, are available and can be queried from existing databases. Where necessary, individuals can be prompted for input via a system-generated questionnaire regarding lifestyle factors, experienced symptoms, and their severity. Data cleaning is also performed in this layer to establish the validity and reliability of the health data.

II. Analysis Layer Situation analysis takes place in this layer, and includes both situation detection and situation prediction.

Situation Detection In the context of AF monitoring, situation detection involves determining whether the monitored person currently has AF or not based on sensor observations and personal information such as experienced symptoms and risk factors. As discussed in Section 2, the detection of AF from sensor observations has been greatly accelerated by ML, with many models meeting and even exceeding human accuracy. In previous work [36, 35], we have implemented ML algorithms such as multilayer perceptron, gradient boosting, and support vector machines to detect AF from ECG data. ML-enabled classifications are captured in the AF moni-

[2] https://www.polar.com/en/sensors/h10-heart-rate-sensor

[3] https://bleak.readthedocs.io

toring ontology (AFMO) as detected situations. This can be used as a screening tool for further clinician-led diagnosis. As shown in Fig. 2, false positives are logged in the system and included as historical training data to improve the ML models. The individual's symptom state (i.e. whether they are symptomatic or asymptomatic) and severity of any symptoms is also captured in the AFMO.

Situation Prediction In this work, we focus primarily on the prediction of AF and its associated health risks, with the two most significant being stroke and major bleeding. AF increases the risk of stroke due to a reduced quality in heart contractions, resulting in slow flow of blood and subsequent formation of blood clots [14]. This is commonly managed using oral anticoagulation to prevent blood clot formation, which in turn has the risk of major bleeding. Therefore, the risks of both stroke and major bleeding must be carefully weighed for people diagnosed with AF.

Risk scores provide a systematic and quantifiable assessment of the likelihood of health outcomes, forming a well-established basis for health situation prediction based on risk factors. We selected the HARMS$_2$-AF score [31], the CHA$_2$DS$_2$-VASc score [21], and the HAS-BLED score [26] to quantify the risks of new onset AF, stroke, and major bleeding respectively. Each scoring system recommends a risk category (low, moderate, or high) based on the score. Tables 2 and 3 summarise the scoring systems and their risk factors respectively.

Table 2: Summary of the scoring systems, where x = risk score.

Scoring System	Risk	Highest Possible Score	Risk Stratification		
			Low Risk	Moderate Risk	High Risk
HARMS$_2$-AF	AF	14	$x \leq 4$	$5 \leq x \leq 9$	$x > 9$
CHA$_2$DS$_2$-VASc	Stroke	10	$x = 0$	$x = 1$	$x \geq 2$
HAS-BLED	Major bleeding	10	$x \leq 1$	$x = 2$	$x \geq 3$

Table 3: Summary of the risk factors for each scoring system.

Scoring System	Risk Factors
HARMS$_2$-AF	**H**ypertension; **A**ge; **R**aised BMI; **M**ale sex; **S**leep apnoea; **S**moking; **A**lcohol
CHA$_2$DS$_2$-VASc	**C**ongestive heart failure/LV dysfunction; **H**ypertension; **A**ge $\geq$ 75; **D**iabetes mellitus; **S**troke; **V**ascular disease; **A**ge btwn. 65 and 74; **S**ex **c**ategory (female)
HAS-BLED	**H**ypertension; **A**bnormal renal/liver function; **S**troke; **B**leeding history/pre-disposition; **L**abile intl. normalized ratio; **E**lderly; **D**rugs/alcohol concomitantly

Since health risk cannot be binary and must take into account the intrinsic uncertainty associated with future outcomes, we propose a fuzzy inference approach to account for uncertainty in the risk scores and health risk categories. The fuzzy inference process for each risk score takes place in three steps:

i. *Fuzzification:* We fuzzify the crisp risk scores using three membership functions for each score corresponding to low, medium, and high score categories.

These input membership functions are based on the recommended risk category stratification for each scoring system, and are shown in Fig. 3. The output of this step is a fuzzified value showing varying degrees of membership to each score category on a scale of 0 to 1. For example, a HARMS$_2$-AF score of 9 has a membership of 0.00 to the low score category, 0.73 to the medium score category, and 0.27 to the high score category. This can be coded as a fuzzy value [0.00, 0.73, 0.27].

ii. *Inference:* To determine the corresponding health risk category, fuzzy rules are used to define the relationship between the input and output. The rules are as follows: if the score is low, then the health risk is low; if the score is medium, then the health risk is moderate; and if the score is high, then the health risk is high. In this case, the score and health risk categories are similar. Therefore, the output of this step will be the same as the input, i.e. [0.00, 0.73, 0.27].

iii. *Defuzzification:* To defuzzify the fuzzy value, a crisp value is computed using Mamdani inference. This value represents the percentage risk. To assign a risk category to the percentage risk, we use another membership function as follows: where y is a person's computed percentage risk, they are at low risk if $0 < y < 40$, at moderate risk if $10 < y < 70$ and at high risk if $40 < y < 100$. This output membership function is shown in Fig. 3. Using this defuzzification method, a HARMS$_2$-AF score of 9 results in a percentage risk of 50.06%, which can be considered mostly moderate risk but also partially high risk as it has a membership of 0.00, 0.68, and 0.32 to the low, moderate, and high risk categories respectively. In contrast, using the recommended HARMS$_2$-AF crisp thresholds, the score of 9 would be considered only moderate risk. Thus, the fuzzy risk category stratification informs appropriate decision support that takes into account the range of risk in each score.

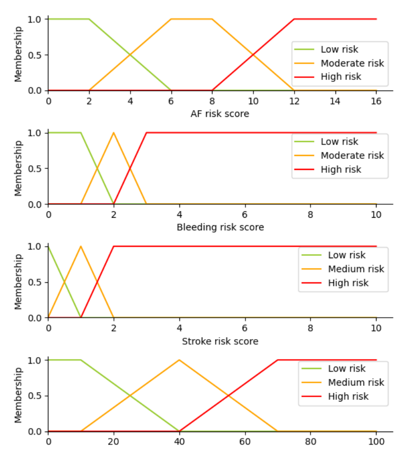

Fig. 3 Membership functions for the HARMS$_2$-AF score, CHA$_2$DS$_2$-VASc score, HAS-BLED score, and the risk category.

III. Decision Layer After situation analysis, it is important that any adverse detected or predicted situations are appropriately managed. Warning alerts are triggered when AF is detected or when moderate or high health risks are predicted. At the same time, recommendations are given based on these adverse situations. The degrees of membership for the health risk categories as determined by fuzzy inference play an important role in the selection of appropriate recommendations. For example, low risk is only considered truly low risk if there is no degree of membership to the high or moderate risk categories. The process flow diagram in Fig. 2 shows the recommendations generated in response to each situation.

5 AF Monitoring Ontology

In this section, we provide details on the development of the AFMO, which provides the data schema for the knowledge graph. Methontology [10] was selected as the ontology development methodology as it is detailed and application-independent. The methodology consists of seven phases which we detail below.

I. Specification In this phase, the AFMO's purpose, scope, and requirements are clearly defined. The main purpose of the ontology within the architecture is to provide representation support and a data schema for three categories of data: sensors and sensor observations, personal information, and expert medical knowledge. This enables higher-level reasoning for situation analysis and decision support. The scope of the AFMO is limited to AF; however, the top-level concepts of the ontology are generalisable for other health use cases. The AFMO requirements are specified in terms of CQs, which cover the representation support and reasoning capabilities of the ontology. Examples of these CQs include: *What are the symptoms of AF? What are the risk factors for AF? What is a particular person's risk level for AF?*

II. Knowledge Acquisition The AFMO was developed using expert knowledge obtained from the literature. Scientific publications, textbooks, and clinical practice guidelines for AF management were consulted.

III. Conceptualization This phase involves structuring three categories of data in a conceptual model. The first category includes classes to represent sensor devices, their measurements, and the properties they measure. The second category encompasses personal information, includes anthropometric, demographic, and lifestyle data, as well as information about the individual's medical history, current diagnoses, symptoms being experienced, and detected and predicted situations. Finally, the medical knowledge category represents risk factors for and symptoms of AF and its associated conditions. It also includes recommendations for mitigating detected and predicted situations. The top level concepts of the AFMO are shown in Fig. 4.

IV. Integration In this phase, the AFMO is integrated with concepts from existing ontologies and vocabularies. Additionally, all health domain concepts are cross-

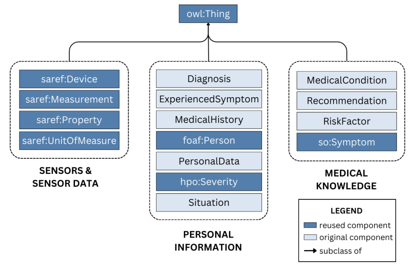

Fig. 4 The top level concepts of the AF monitoring ontology.

referenced to the Unified Medical Language System (UMLS) metathesaurus [3], which maps to several medical vocabularies and standards.

V. Implementation In this phase, the AFMO is codified in a formal language. The AFMO was implemented using Protégé[4]. Rules and queries were written using Semantic Web Rule Language (SWRL)[5] and SPARQL Protocol and RDF Query Language (SPARQL)[6] respectively.

VI. Evaluation Throughout the ontology development process, the Pellet reasoner was used to detect inconsistencies in the AFMO. No inconsistencies were detected in the final version of the AFMO. Additionally, the OntOlogy Pitfall Scanner! (OOPS!) [27] was used to evaluate different aspects of the AFMO including modelling decisions and inferences. Only one minor pitfall was detected: “P2: Using different naming conventions in the ontology”. This is because of the re-use of concepts from existing ontologies which all have different naming conventions. Finally, the answerability of the CQs was also evaluated using queries.

VII. Documentation The AFMO is publicly available at a persistent URL (PURL)[7]. The AFMO specification and documentation was created using the WIzard for DOCumenting Ontologies (WIDOCO) [12] and is publicly accessible online[8].

6 Evaluation and Analysis

Use Case Evaluation To evaluate the representation support and reasoning capability of the architecture for the use case, we created 25 synthetic user profiles with various random combinations of medical histories, diagnoses, symptoms, and other personal information. Because the primary focus of this work is on risk prediction using risk scores, no sensor data was included in the synthetic profiles. The profiles’ ages range from 31 to 80, with 11 being male and 14 being female, and 12 having an

[4] https://protege.stanford.edu/
[5] https://www.w3.org/Submission/SWRL/
[6] https://www.w3.org/TR/rdf-sparql-query/
[7] https://purl.org/afmo
[8] https://mbithenzomo.github.io/afmo

existing diagnosis of AF. As per the process flow diagram in Fig. 2, the HARMS$_2$-AF score is computed for those who do not have an AF diagnosis, while the CHA$_2$DS$_2$-VASc and HAS-BLED scores are computed for those with an AF diagnosis. Table 4 gives a partial summary of the synthetic user profiles. The full synthetic data as well as the code used to generate it is available on GitHub[9].

Table 4: Partial summary of the synthetic user profiles. M: male; F: female; Y: yes; N: no; LR: low risk; MR: moderate risk; HR: high risk; N/A: not applicable

ID	Age	Sex	BMI	Smoking Status	Weekly Alcoholic Drinks	AF Diagnosis	Crisp Risk Score and Category		
							AF	Stroke	Major Bleeding
001	50	M	28.1	Never	13	Y	N/A	1 (MR)	3 (HR)
002	47	F	24.8	Former	14	N	1 (LR)	N/A	N/A
003	64	M	26.5	Former	4	N	7 (MR)	N/A	N/A
004	43	F	19.5	Current	9	Y	N/A	5 (HR)	7 (HR)
005	33	M	31.6	Former	4	Y	N/A	4 (HR)	2 (MR)
006	49	F	38.0	Never	16	Y	N/A	3 (HR)	3 (HR)
007	72	F	31.0	Former	11	N	4 (LR)	N/A	N/A
008	32	F	29.0	Never	14	N	5 (MR)	N/A	N/A
009	75	F	28.9	Never	7	Y	N/A	6 (HR)	5 (HR)
010	43	F	30.3	Former	4	N	1 (LR)	N/A	N/A
011	68	M	29.0	Current	15	N	9 (MR)	N/A	N/A
012	67	F	29.6	Never	10	N	3 (LR)	N/A	N/A
013	51	M	22.1	Never	2	N	2 (LR)	N/A	N/A
014	59	F	24.6	Former	7	Y	N/A	5 (HR)	4 (HR)
015	62	M	25.3	Never	5	Y	N/A	3 (HR)	3 (HR)
016	60	F	28.4	Current	1	N	6 (MR)	N/A	N/A
017	46	M	27.2	Current	10	N	6 (MR)	N/A	N/A
018	62	M	27.9	Current	12	Y	N/A	1 (MR)	2 (MR)
019	54	F	43.7	Never	13	Y	N/A	2 (HR)	5 (HR)
020	31	M	22.2	Never	11	Y	N/A	2 (HR)	3 (HR)
021	70	F	30.1	Former	8	Y	N/A	5 (HR)	5 (HR)
022	44	M	23.9	Never	11	N	3 (LR)	N/A	N/A
023	47	F	22.2	Former	9	N	1 (LR)	N/A	N/A
024	80	F	36.7	Never	17	N	7 (MR)	N/A	N/A
025	78	M	24.9	Current	4	Y	N/A	3 (HR)	4 (HR)

After creating the synthetic profiles, we ran the Pellet reasoner to confirm that the correct inferences were computed, and queried the knowledge graph to answer the CQs. For example, Fig. 5(a) shows a subset of the property assertions inferred by the Pellet reasoner for a particular profile, including predicted bleeding and stroke risk levels as well as the corresponding recommendations, while Fig. 5(b) shows a sample query to list all individuals at high risk for stroke. The properties asserted by the reasoner and the results of the sample query show that the architecture is able to correctly categorise the risk of AF and its associated conditions, as well as give appropriate recommendations to mitigate these risks.

[9] https://github.com/mbithenzomo/afmo

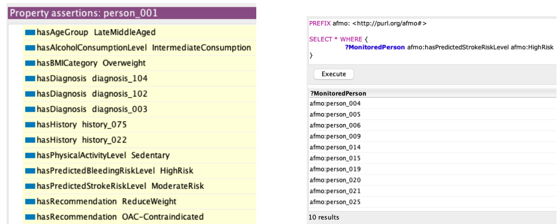

(a) Property assertions (b) Sample SPARQL query

Fig. 5 (a) A subset of the property assertions inferred by the Pellet reasoner. (b) A sample SPARQL query showing individuals with high risk levels for stroke.

Architectural Analysis Having demonstrated the implementation of the architecture using the AF use case, we now analyse it according to the previously defined design goals.

- CQ1 Process/clinical interoperability: This is a pivotal thread running through the entire architecture and monitoring process, supported by incorporating established clinical practice guidelines from accredited health organisations.
- CQ2 Semantic interoperability: The ontology supports the cross-referencing of health concepts to standard medical terminologies through the UMLS.
- CQ3 Syntactic interoperability: This is achieved in the data layer, in which all data and resources are mapped to the FHIR standard.
- CQ4 Contextual information: This is supported through the ontology, which facilitates the representation of contextual data including identity (e.g. name and age) and temporal concepts (e.g. timestamps and GPS coordinates).
- CQ5 Situation analysis: The architecture supports both rule-based situation prediction using health risk scores and the capture of ML classifications in the ontology for situation detection.
- CQ6 Transparency and explainability: Semantic web technologies, which are highly interpretable and enhance explainability in AI systems, are at the core of the architecture. Additionally, rule-based reasoning contributes to transparency in situation analysis and decision support.
- CQ7 Recommendations and warnings: The decision layer supports the generation of warnings triggered by adverse situations, and recommendations to mitigate these situations based on clinical guidelines.
- CQ8 Data pre-processing: This is done in the data layer, where signal processing and data cleaning techniques are implemented to ensure the data is valid and suitable for semantic annotation and situation analysis.
- CQ9 Uncertainty handling: The architecture supports fuzzy inference in the analysis layer, allowing for varying degrees of membership to health categories as part of situation analysis, which in turn influences decision support.
- CQ10 Representation support: This is facilitated by the ontology, which forms the core of the architecture. All health data and expert knowledge is semantically annotated and captured in the ontology.

7 Discussion and Conclusion

This paper proposes a semantic architecture for personal health monitoring using heterogeneous sources of health data. Through an ontology, the architecture can semantically represent health data and expert knowledge. Rules and fuzzy inference are used for reasoning, thereby facilitating situation analysis and decision support under uncertainty. The functionality of the architecture is demonstrated through a prototype implementation for the use case of AF. We develop an ontology, AFMO, to model concepts relating to sensor observations, personal information, and medical knowledge about AF. These core concepts are generalisable across multiple health monitoring applications. A knowledge graph is then built from the ontology using synthetic user profiles. Through the use case, we demonstrate that the architecture can support risk detection, prediction, and decision making.

As demonstrated by experiments run on the synthetic user profiles, the architecture can correctly categorise the risk of AF and its associated conditions, and give appropriate recommendations to mitigate these risks. We demonstrate how the recommendations maintain human involvement, supporting clinician-led, individual-led, and shared interventions. Importantly, we show that the monitoring process is grounded in established clinical guidelines, demonstrating the architecture's capability to support process and clinical interoperability. Additionally, we demonstrate that the architecture addresses key challenges identified in health monitoring systems.

There are some limitations in the proposed approach. Firstly, rules are time-consuming to develop and require updates to remain adaptive. This can be mitigated using ML to automate the creation of dynamic rules, although these may need to undergo verification from domain experts. Secondly, although fuzzy inference provides a framework for representing the degrees of membership to defined categories, it is limited in its ability to explain the uncertainty in causality, which is beneficial for explaining the impact of risk factors on health. Additionally, the risk scores assume certainty in the inputted risk factors. However, these inputs may themselves be uncertain since they rely on precise and comprehensive health records which may not be available. Probabilistic models such as BNs are well suited for representing uncertainty in causal risk factors, and the combination of ontologies with BNs has seen some success. Subsequent iterations of the architecture will implement both fuzzy inference and BNs, and be evaluated on other health monitoring applications and with real-world data to demonstrate its generalisability.

Finally, while continuous monitoring is a promising solution to the increasing prevalence of non-communicable diseases, it is not without its risks. Although false positives are preferred to false negatives, they may cause undue anxiety and psychological harm to the monitored person and result in costly and/or invasive diagnostic procedures. These potential harms must be carefully considered by both monitored persons and clinicians before the adoption of health monitoring systems.

Acknowledgements This work was financially supported by the Hasso Plattner Institute for Digital Engineering through the HPI Research School at UCT. It was also supported in part by the National Research Foundation of South Africa (Grant Numbers: 151217, SRUG2204264808).

References

1. Ali, F., Islam, S.M.R., Kwak, D., Khan, P., Ullah, N., Yoo, S.j., Kwak, K.S.: Type-2 fuzzy ontology–aided recommendation systems for IoT–based healthcare. Computer Communications **119**, 138–155 (2018)
2. Benson, T., Grieve, G.: Principles of Health Interoperability: FHIR, HL7 and SNOMED CT. Health Information Technology Standards, Springer International Publishing (2021)
3. Bodenreider, O.: The Unified Medical Language System (UMLS): integrating biomedical terminology. Nucleic Acids Research **32**(Database issue), D267–D270 (2004)
4. Brieger, D., Amerena, J., Attia, J., Bajorek, B., Chan, K.H., Connell, C., et al.: Australian Clinical Guidelines for the Diagnosis and Management of Atrial Fibrillation 2018. Heart, Lung and Circulation **27**(10), 1209–1266 (2018)
5. Burnham, J.P., Lu, C., Yaeger, L.H., Bailey, T.C., Kollef, M.H.: Using wearable technology to predict health outcomes: a literature review. Journal of the American Medical Informatics Association **25**(9), 1221–1227 (2018)
6. De Brouwer, M., Steenwinckel, B., Fang, Z., Stojchevska, M., Bonte, P., De Turck, F., et al.: Context-aware query derivation for IoT data streams with DIVIDE enabling privacy by design. Semantic Web **14**(5), 893–941 (2023)
7. Dunn, J., Kidzinski, L., Runge, R., Witt, D., Hicks, J.L., Schüssler-Fiorenza Rose, S.M., et al.: Wearable sensors enable personalized predictions of clinical laboratory measurements. Nature Medicine **27**(6), 1105–1112 (2021)
8. Esposito, M., Minutolo, A., Megna, R., Forastiere, M., Magliulo, M., De Pietro, G.: A smart mobile, self-configuring, context-aware architecture for personal health monitoring. Engineering Applications of Artificial Intelligence **67**, 136–156 (2018)
9. Feldman, K., Duncan, R.G., Nguyen, A., Cook-Wiens, G., Elad, Y., Nuckols, T., Pevnick, J.M.: Will Apple devices' passive atrial fibrillation detection prevent strokes? Estimating the proportion of high-risk actionable patients with real-world user data. Journal of the American Medical Informatics Association **29**(6), 1040–1049 (2022)
10. Fernández-López, M., Gómez-Pérez, A., Juristo, N.: METHONTOLOGY: From Ontological Art Towards Ontological Engineering. In: Proceedings of the Ontological Engineering AAAI-97 Spring Symposium Series. pp. 33–40 (1997)
11. Gambhir, S.S., Ge, T.J., Vermesh, O., Spitler, R., Gold, G.E.: Continuous health monitoring: An opportunity for precision health. Science Translational Medicine **13**(597), eabe5383 (2021)
12. Garijo, D.: Widoco: a wizard for documenting ontologies. In: International Semantic Web Conference. pp. 94–102 (2017)
13. Hagras, H.: Toward Human-Understandable, Explainable AI. Computer **51**(9), 28–36 (2018)
14. Hindricks, G., Potpara, T., Dagres, N., Arbelo, E., Bax, J.J., Blomström-Lundqvist, C., et al.: 2020 ESC Guidelines for the diagnosis and management of atrial fibrillation developed in collaboration with the European Association for Cardio-Thoracic Surgery (EACTS). European Heart Journal **42**(5), 373–498 (2021)
15. Hitzler, P.: Semantic Web: A Review Of The Field. Communications of the ACM **64**(2), 76–83 (2021)
16. Hong, S., Zhou, Y., Shang, J., Xiao, C., Sun, J.: Opportunities and challenges of deep learning methods for electrocardiogram data: A systematic review. Computers in Biology and Medicine **122**, 103801 (2020)
17. Hristoskova, A., Sakkalis, V., Zacharioudakis, G., Tsiknakis, M., De Turck, F.: Ontology-Driven Monitoring of Patient's Vital Signs Enabling Personalized Medical Detection and Alert. Sensors **14**(1), 1598–1628 (2014)
18. Huang, J.D., Wang, J., Ramsey, E., Leavey, G., Chico, T.J.A., Condell, J.: Applying Artificial Intelligence to Wearable Sensor Data to Diagnose and Predict Cardiovascular Disease: A Review. Sensors **22**(20), 8002 (2022)
19. Kordestani, H., Mojarad, R., Chibani, A., Barkaoui, K., Amirat, Y., Zahran, W.: Extended Hapicare: A telecare system with probabilistic diagnosis and self-adaptive treatment. Expert Systems with Applications **186**, 115749 (2021)

20. Kornej, J., Börschel, C.S., Benjamin, E.J., Schnabel, R.B.: Epidemiology of Atrial Fibrillation in the 21st Century Novel Methods and New Insights. Circulation Research **127**, 4–20 (2020)
21. Lip, G.Y., Nieuwlaat, R., Pisters, R., Lane, D.A., Crijns, H.J.: Refining Clinical Risk Stratification for Predicting Stroke and Thromboembolism in Atrial Fibrillation Using a Novel Risk Factor-Based Approach. Chest **137**(2), 263–272 (Feb 2010)
22. Mcheick, H., Nasser, H., Dbouk, M., Nasser, A.: Stroke Prediction Context-Aware Health Care System. In: 2016 IEEE International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE). pp. 30–35 (2016)
23. Moodley, D., Seebregts, C.: Re-imagining health and well-being in low resource African settings using an augmented AI system and a 3D digital twin. In: Proceedings of the 1st Workshop on AI for Digital Twins and Cyber-Physical Applications, in conjunction with IJCAI 2023. vol. 3541 (2023)
24. Nzomo, M., Moodley, D.: Semantic Technologies in Sensor-Based Personal Health Monitoring Systems: A Systematic Mapping Study. arXiv preprint arXiv:2306.04335 (2023)
25. Pereira, T., Tran, N., Gadhoumi, K., Pelter, M.M., Do, D.H., Lee, R.J., Colorado, R., Meisel, K., Hu, X.: Photoplethysmography based atrial fibrillation detection: a review. npj Digital Medicine **3**(1), 3 (2020)
26. Pisters, R., Lane, D.A., Nieuwlaat, R., de Vos, C.B., Crijns, H.J.G.M., Lip, G.Y.H.: A Novel User-Friendly Score (HAS-BLED) To Assess 1-Year Risk of Major Bleeding in Patients With Atrial Fibrillation: The Euro Heart Survey. Chest **138**(5), 1093–1100 (2010)
27. Poveda-Villalón, M., Gómez-Pérez, A., Suárez-Figueroa, M.C.: OOPS! (OntOlogy Pitfall Scanner!): An On-line Tool for Ontology Evaluation. International Journal on Semantic Web and Information Systems **10**(2), 7–34 (2014)
28. Ramella-Roman, J.C.: 22. Sources of inaccuracy in wearable photoplethysmography. In: The 2023 Wearable Photoplethysmography Roadmap, pp. 108–110. Physiological Measurement (2023)
29. Ren, Y., Parvizi, A., Mellish, C., Pan, J.Z., Van Deemter, K., Stevens, R.: Towards Competency Question-Driven Ontology Authoring. In: Proceedings of the European Semantic Web Conference (ESWC 2014). vol. 8465, pp. 752–767 (2014)
30. Saarinen, H.J., Joutsen, A., Korpi, K., Halkola, T., Nurmi, M., Hernesniemi, J., Vehkaoja, A.: Wrist-worn device combining PPG and ECG can be reliably used for atrial fibrillation detection in an outpatient setting. Frontiers in Cardiovascular Medicine **10** (2023)
31. Segan, L., Canovas, R., Nanayakkara, S., Chieng, D., Prabhu, S., Voskoboinik, A., et al.: New-onset atrial fibrillation prediction: the HARMS2-AF risk score. European Heart Journal p. ehad375 (2023)
32. Serhal, H., Abdallah, N., Marion, J.M., Chauvet, P., Oueidat, M., Humeau-Heurtier, A.: Overview on prediction, detection, and classification of atrial fibrillation using wavelets and AI on ECG. Computers in Biology and Medicine **142**, 105168 (2022)
33. Spertus, J., Dorian, P., Bubien, R., Lewis, S., Godejohn, D., Reynolds, M.R., et al.: Development and validation of the AFEQT Questionnaire in patients with atrial fibrillation. Circulation: Arrhythmia and Electrophysiology **4**(1), 15–25 (2011)
34. Tamašauskaitė, G., Groth, P.: Defining a Knowledge Graph Development Process Through a Systematic Review. ACM Transactions on Software Engineering and Methodology **32**(1), 1–40 (2023)
35. Wanyana, T., Nzomo, M., Price, C.S., Moodley, D.: Combining Machine Learning and Bayesian Networks for ECG Interpretation and Explanation. In: Proceedings of the 8th International Conference on Information and Communication Technologies for Ageing Well and e-Health (ICT4AWE 2022). pp. 81–92 (2022)
36. Wanyana, T., Nzomo, M., Price, C.S., Moodley, D.: A Personal Health Agent for Decision Support in Arrhythmia Diagnosis. In: Maciaszek, L.A., Mulvenna, M.D., Ziefle, M. (eds.) Information and Communication Technologies for Ageing Well and e-Health. pp. 385–407. Springer Nature Switzerland, Cham (2023)