

A probabilistic ontology-based platform for self-learning context-aware healthcare applications

Femke Ongenae^{a,*}, Maxim Claeys^a, Thomas Dupont^a, Wannes Kerckhove^a,
Piet Verhoeve^b, Tom Dhaene^a, Filip De Turck^a

^a*Department of Information Technology (INTEC), Ghent University - iMinds, Gaston Crommenlaan 8 bus 201, B-9050 Ghent, Belgium*

^b*iMinds VZW, Gaston Crommenlaan 8 bus 102, B-9050 Ghent, Belgium*

Abstract

Context-aware platforms consist of dynamic algorithms that take the context information into account to adapt the behavior of the applications. The relevant context information is modeled in a context model. Recently, a trend has emerged towards capturing the context in an ontology, which formally models the concepts within a certain domain, their relations and properties.

Although much research has been done on the subject, the adoption of context-aware services in healthcare is lagging behind what could be expected. The main complaint made by users is that they had to significantly alter workflow patterns to accommodate the system. When new technology is introduced, the behavior of the users changes to adapt to it. Moreover, small differences in user requirements often occur between different environments where the application is deployed. However, it is difficult to foresee these

*Corresponding author: Tel.: +32 9 331 49 38, Fax: +32 9 331 48 99

Email addresses: Femke.Ongenae@intec.UGent.be (Femke Ongenae), Maxim.Claeys@intec.UGent.be (Maxim Claeys), Thomas.Dupont@intec.UGent.be (Thomas Dupont), Wannes.Kerckhove@intec.UGent.be (Wannes Kerckhove), Piet.Verhoeve@iminds.be (Piet Verhoeve), Tom.Dhaene@intec.UGent.be (Tom Dhaene), Filip.DeTurck@intec.UGent.be (Filip De Turck)

changes in workflow patterns and requirements at development time. Consequently, the context-aware applications are not tuned towards the needs of the users and they are required to change their behavior to accommodate the technology instead of the other way around.

To tackle this issue, a self-learning, probabilistic, ontology-based framework is proposed, which allows context-aware applications to adapt their behavior at run-time. It exploits the context information gathered in the ontology to mine for trends and patterns in the behavior of the users. These trends are then prioritized and filtered by associating probabilities, which express their reliability. This new knowledge and their associated probabilities are then integrated into the context model and dynamic algorithms. Finally, the probabilities are increased or decreased, according to context and behavioral information gathered about the usage of the learned information.

A use case is presented to illustrate the applicability of the framework, namely mining the reasons for patients' nurse call light use to automatically launch calls. Detecting Systemic Inflammatory Response Syndrome (SIRS) as a reason for nurse calls is used as a realistic scenario to evaluate the correctness and performance of the proposed framework. It is shown that correct results are achieved when the dataset contains at least 1,000 instances and the amount of noise is lower than 5%. The execution time and memory usage are also negligible for a realistic dataset, i.e., below 100 ms and 10 MB.

Keywords:

Context-aware, Self-learning, Ontology, Probability, eHealth, Nurse call system

1. Introduction

Computerized tools, health monitoring devices and sensors are being actively adopted in modern healthcare settings, especially to support administrative tasks, data management and patient monitoring (Orwat et al., 2008; Colpaert et al., 2009). Today, caregivers are directly faced with these technologies, which increases the complexity of their daily activities (Tentori et al., 2009). The caregiver has to use several devices to manually consult, insert and combine data, even when carrying out a single task. This is very time-consuming. Due to this inadequate integration of the technology, as well as the large amount of data being generated by the devices and the heavy workload of staff members, it is not rare for important events to be missed, e.g., early indications of worsening condition of a patient. To resolve this issue, context-aware techniques are often proposed to automatically exploit the medical information available to improve continuous care and personalize healthcare (Burgelman and Punie, 2006).

Although much research has been done on the subject, the adoption of context-aware services is lagging behind what could be expected. Most of the projects are prototypes and real applications are still difficult to find. Whereas the healthcare industry is quick to exploit the latest medical technology, they are reluctant adopters of modern health information systems (Chin, 2004). Half of all computer-based information systems fail due to user resistance and staff interference (Anderston and Aydin, 1997). The main complaint made against mobile, context-aware systems is that users had to significantly alter workflow patterns to accommodate the system (Jahnke et al., 2004). This is due to inadequate techniques for personalization of the ser-

26 vices, a lack of focus on the soft aspects of interaction, e.g., automated and
27 personalized alerts, and the lack of tackling problems such as the need of the
28 users for control (Criel and Claeys, 2008).

29 The context-aware platforms use dynamic algorithms, which take the con-
30 text information into account, to adapt the behavior of the applications ac-
31 cording to the context and offer personalized services to the users. However,
32 these algorithms are defined at development time. When new technology
33 is introduced, the behavior of the users changes to adapt to it. Moreover,
34 different environments in which the application is deployed, e.g., different
35 nursing units or hospital departments, might have slightly different require-
36 ments pertaining to how the context information is taken into account. It is
37 difficult to foresee these changes in behavior and small nuances in workflows
38 at development time. This means that the context model might be incom-
39 plete or the algorithms of the applications built on it may no longer apply.
40 As the applications do not adapt to the requirements and workflow patterns
41 of the users, they feel less in control of the technology and have to adapt their
42 behavior to accommodate the technology instead of the other way around.

43 To tackle this issue, this paper proposes a self-learning framework, which
44 allows the context-aware applications to adapt their behavior at run-time to
45 accommodate the changing requirements of the users. The proposed frame-
46 work consist of the following techniques. First, an ontology-based context
47 model with accompanying rule-based context-aware algorithms is used to
48 capture the behavior of the user and the context in which it is exhibited.
49 This captured information is then filtered, cleaned and structured so that it
50 can be used as input for data mining techniques. The results of these data

51 mining techniques are then prioritized and filtered by associating probabil-
52 ities with the obtained results expressing how reliable or accurate they are.
53 These results and their associated probabilities are then integrated into the
54 context model and dynamic algorithms. These probabilities clarify to the
55 stakeholders that this new knowledge has not been confirmed by rigorous
56 evaluation. Finally, the probabilities are adapted, i.e., in- or decreased, ac-
57 cording to context and behavioral information gathered about the usage of
58 the learned information.

59 The remainder of this article is organized as follows. In Section 2 the
60 relevant related work is discussed and our contribution is highlighted. Sec-
61 tion 3 presents the architecture of the proposed probabilistic ontology-based
62 framework for self-learning context-aware healthcare applications. Section 4
63 discusses the generic implementation of the framework, i.e., the classes that
64 can be extended to implement the specific use cases. The implementation
65 of a specific use case, namely mining the reasons for patients’ call light use
66 to automatically launch calls, is presented in Section 5. Finally, the main
67 conclusions of this research are highlighted and the future work is discussed
68 in Section 6.

69 **2. Related work**

70 *2.1. Context-aware systems*

71 Dey and Abowd (2000) refer to context as “any information that can be
72 used to characterize the situation of entities (i.e., whether a person, place or
73 object) that are considered relevant to the interaction between a user and an
74 application, including the user and the application themselves”. A system

75 may be labeled as “context-aware” if it can acquire, interpret and use context
 76 information to adapt its behavior to the current context in use (Byun and
 77 Cheverst, 2004). A number of generic context platforms have been developed
 78 to relieve application developers from the aggregation and abstraction of con-
 79 text information and the derivation of high-level contexts (Hong et al., 2009a;
 80 Baldauf et al., 2007; Xue and Pung, 2012; Yilmaz and Erdur, 2012). Unor-
 81 ganized, unprocessed raw data can be voluminous, but has no meaning on
 82 itself as it has no relationships or context. Information is data that has been
 83 given meaning by defining relational connections. The proposed platforms
 84 employ several techniques to model this context information, i.e., key-value,
 85 markup scheme, graphical, object-oriented, logic-based and ontology-based
 86 models (Strang and Linnhoff-Popien, 2004). A notable trend is emerging to-
 87 wards ontology-based context-aware platforms (Gu et al., 2005; Chen, 2004;
 88 Santos et al., 2007; Román et al., 2002).

89 To write the dynamic algorithms, which take the context information
 90 captured in the ontology into account to achieve personalized and context-
 91 aware applications, two approaches are commonly used, namely rules or ma-
 92 chine learning techniques (Tsang and Clarke, 2008). Rules are manually con-
 93 structed at development time and thus require developers to foresee all pos-
 94 sible situations that can occur at runtime and define the appropriate corre-
 95 sponding actions. Rules are difficult to modify, maintain and scale (Prentzas
 96 and Hatzilygeroudis, 2007). Machine learning techniques, e.g., Bayesian net-
 97 works and neural networks, are also trained at development time. Bayesian
 98 networks suffer from similar maintenance and scalability problems as the rule-
 99 based approach and acquiring accurate probabilities is a tedious job (Russell

100 and Norvig, 2003). Neural Networks require a lot of processing power and
101 have consequently only been sparsely applied in context-aware applications.
102 Their black-box nature also makes it difficult to gain insight into relations
103 between context and actions, increasing the fear of technology and loss of
104 control from the users. Consequently, with each of these approaches, the
105 context-aware system is only able to cope with a fixed set of context changes
106 that were taken into account during the design of the system.

107 As mentioned previously, run-time adaptation of the dynamic algorithms
108 is needed to adapt to changing behavior of the stakeholders and to truly offer
109 personalized services tuned to the work practices of the specific environment
110 where the application is deployed. A couple of context-aware systems exist
111 that try to tackle this problem by mining historical information (Tsang and
112 Clarke, 2008; Baralis et al., 2011; Strobbe et al., 2012a; Hong et al., 2009b).
113 However, most of the research focusses on the development of data mining
114 techniques, which can be used to learn the patterns and requirements, or use
115 a black-box approach. Little research has been done on the development of a
116 complete framework for self-learning, context-aware applications and on how
117 the learned knowledge should be integrated in an ontology-based platform.

118 *2.2. Context-aware systems in healthcare*

119 The use of context and context-awareness in healthcare is an active re-
120 search area (Bricon-Souf and Newman, 2007; Varshney, 2009). First, there
121 is a large amount of available information, specific healthcare situations and
122 related tasks, which create a potential for cognitive overload amongst the
123 caregivers. Second, the patients, healthcare professionals and some equip-
124 ment are fairly mobile, which requires accurate localization and adaptation

125 of the healthcare services to the environment. Third, the financial and human
 126 resources are limited. This implies a need to cut cost while improving the
 127 quality of service to an increased number of people. Context-aware and per-
 128 vasive prototypes have been developed for a number of hospital (Bardram,
 129 2004; Skov and Hoegh, 2006; Mitchell et al., 2000; Stanford, 2003; Munoz
 130 et al., 2003) and homecare & residential care (Fishkin et al., 2003; Floerke-
 131 meier and Siegemund, 2003; Korhonen et al., 2003; de Toledo et al., 2006;
 132 Hu et al., 2010; Mihailidis et al., 2003; Suzuki and Doi, 2001; Jansen and
 133 Deklerck, 2006) use cases. Examples of context-aware healthcare systems
 134 based on ontologies can also be found in literature (Fook et al., 2006; Zhang
 135 et al., 2005; Paganelli and Giuli, 2011; Ongenaes et al., 2011d).

136 *2.3. eHealth ontologies*

137 An ontology (Gruber, 1993) is a semantic model that formally describes
 138 the concepts in a certain domain, their relationships and attributes. In this
 139 way, an ontology encourages re-use and integration. By managing the data
 140 about the current context in an ontology, intelligent algorithms that take ad-
 141 vantage of this information to optimize and personalize the context-aware
 142 applications, can more easily be defined. The Web Ontology Language
 143 (OWL) (McGuinness and Harmelen, 2004) is the leading language for en-
 144 coding these ontologies. Because of the foundation of OWL in Description
 145 Logics (DLs) (Baader et al., 2003), which are a family of logics that are de-
 146 cidable fragments of first-order logic, the models and description of data in
 147 these models can be formally proved. It can also be used to detect inconsis-
 148 tencies in the model as well as infer new information out of the correlation of
 149 this data. This proofing and classification process is referred to as Reason-

150 ing. Reasoners are implemented as generic software-modules, independent
151 of the domain-specific problem. Ontologies thus effectively separate the do-
152 main knowledge, which can be re-used across different applications, from the
153 application logic, which can be written as rules on top of the ontology.

154 The definition and use of ontologies in the medical domain is an ac-
155 tive research field, as it has been recognized that ontology-based systems
156 can be used to improve the management of complex health systems (Valls
157 et al., 2010). Most of the developed ontologies focus on biomedical research
158 and are mainly employed to clearly define medical terminology (Ongenaes
159 et al., 2011b), e.g., Galen Common Reference Model (Rector et al., 2003),
160 the Foundational Model of Anatomy Ontology (FMA) (Rosse and Jr, 2008)
161 or the Gene Ontology (Blake and Harris, 2008). Little work has been done
162 on developing high-level ontologies, which can be used to model context
163 information and knowledge utilized across the various continuous care set-
164 tings (Ongenaes et al., 2011a). However, ontologies have been developed for
165 specific subdomains of continuous care, e.g., ontologies for structuring organi-
166 zation knowledge in homecare assistance (Valls et al., 2010), representing the
167 context of the activity in which the user is engaged (Rodríguez et al., 2011)
168 and modeling chronic disease management in homecare settings (Paganelli
169 and Giuli, 2011).

170 *2.4. Our contribution*

171 In this paper, we propose a self-learning and probabilistic framework
172 to adapt the behavior of ontology-based, context-aware applications to the
173 changing requirements of the users and their workflow patterns. To our
174 knowledge, little previous research has been done on how discovered trends

175 and patterns can be integrated into ontology-based platforms without making
176 the existing model inconsistent. To tackle this issue, we use a probabilistic
177 approach, which conveys the reliability of the learned knowledge to the users
178 and ensures the compatibility with existing knowledge in the context model.
179 Moreover, the existing research on self-learning, context-aware applications
180 concentrates on exploring data mining techniques, which can be used to dis-
181 cover the trends and patterns. Our research focuses on the development
182 of a complete framework to enable self-learning, context-aware healthcare
183 applications.

184 **3. Architecture of the self-learning, context-aware framework**

185 The general architecture of the proposed self-learning, context-aware frame-
186 work is visualized in Figure 1. The following subsections discuss the different
187 components and modules of this framework in more detail.

188 *3.1. Context-aware platform*

189 The general architecture of a context-aware, ontology-based platform can
190 be split up into five layers. The *Device Layer* includes all the devices and
191 the software on those devices that deliver context information. The modern
192 healthcare settings contains a plethora of computerized medical equipment
193 to convey the condition of a patient, e.g., monitoring equipment, electronic
194 patient records and laboratory results stored in a database, and support the
195 caregivers in their daily activities, e.g., nurse call systems and task manage-
196 ment and planning tools.

197 The *Context Provider Layer* takes care of the acquisition of specific con-
198 text information, e.g., location or presence information, and translates it to

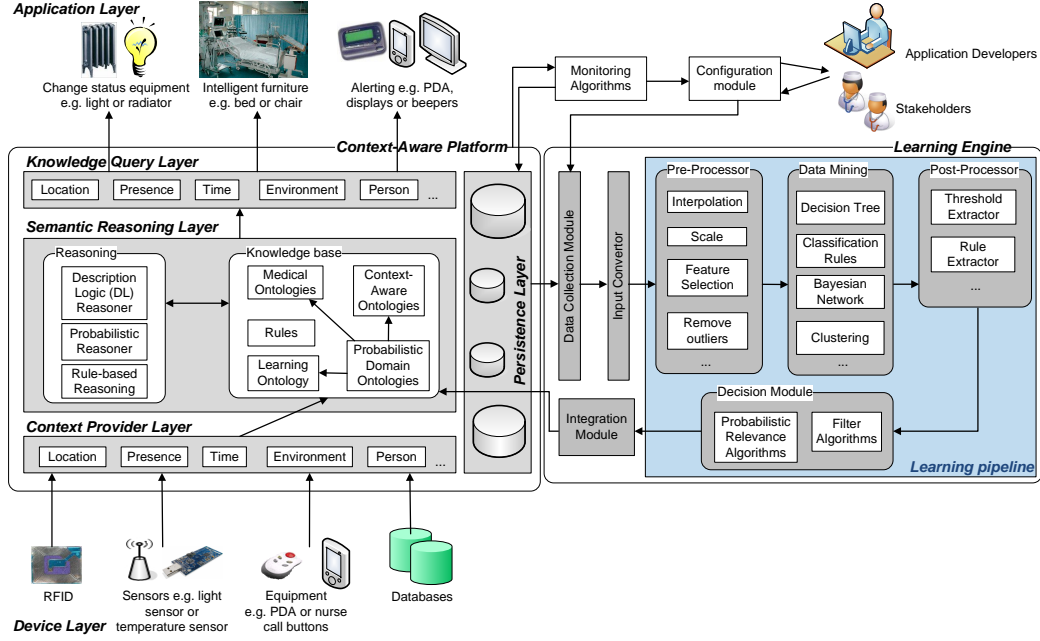


Figure 1: General architecture of the self-learning, context-aware framework

ontology instances. These ontology instances are then added to the *Knowledge Base* in the *Semantic Reasoning Layer*. This *Knowledge Base* aggregates all the relevant context information into a formal context model, i.e., an ontology. Existing *Medical and Context-Aware Ontologies* are integrated into the platform and extended with *Domain Ontologies* which model the information specific to a particular healthcare setting, e.g., the specific roles and competences of the caregivers and how they map on each other, the available monitoring equipment and their threshold values and specific tasks that need to be performed. These *Domain Ontologies* can also contain probabilistic information, e.g., a call made by patient with a heart disease has 25% chance of being urgent.

Reasoning components are then used to derive new, high-level knowledge

211 from the information aggregated in the *Knowledge Base*. Due to the foun-
212 dation of ontologies in *Description Logics (DL)*, the models can be formally
213 proofed by using a *DL Reasoner*. This *DL Reasoner* is used to detect incon-
214 sistencies in the model as well as infer new information from the correlation of
215 the data. For example, a concept *Fever* is created in the ontology, which au-
216 tomatically detects patients with a temperature above 38 °C. More complex
217 logic is expressed by defining *Rules* on top of this ontology and performing
218 *Rule-based Reasoning*.

219 The *Knowledge Query Layer* facilitates the retrieval of context informa-
220 tion such that it can be used by the different applications and services. The
221 *Application Layer* includes all the devices and the software on those devices
222 that use the (derived) context information to adapt their behavior.

223 Finally, the *Persistence Layer* ensures the persistence of context informa-
224 tion. Static contextual information about users, devices and the environment
225 can be easily obtained from these databases. More importantly, the *Persis-*
226 *tence Layer* can also be used to store more dynamic information, such as
227 previous locations of caregivers and patients or actions taken by the users.

228 The *Semantic Reasoning* and *Persistence Layers* are the most important
229 layers to facilitate a self-learning *Context-Aware Platform*. As the *Knowledge*
230 *Base* integrates all the context information, it gives insight into the behavior
231 and changing requirements of the users. All the collected context information
232 and the knowledge derived from it is then persisted in the databases from
233 the *Persistence Layer*. This lets the *Learning Engine* exploit this history of
234 context information to derive trends and patterns and adapt the information
235 in the ontology and accompanying rules accordingly.

236 3.2. Monitoring algorithms and configuration module

237 *Monitoring Algorithms* determine missing or inaccurate knowledge in the
238 ontology. An example: situations are logged where a suggestion is given by
239 the system to the staff to do an action, but under certain circumstances the
240 caregivers consistently execute a different action. The *Monitoring Algorithms*
241 constantly monitor the ontology for interesting situations. They gather these
242 situations and store them collectively in the *Persistence Layer*. The results
243 of the *Monitoring Algorithms* can intermediately be shown to *Stakeholders*,
244 i.e., domain experts such as nurses, doctors and professionals working for the
245 healthcare industry, and *Application Developers*. When enough data has been
246 collected, the *Learning Engine* can be initiated. The amount of data that
247 should be gathered depends on the specific use case and the used data mining
248 technique. The input parameters are specified in the *Configuration Module*
249 and the *Data Collection Module* automatically extracts the appropriate data
250 from the *Persistence Layer*. The *Configuration Module* is also responsible
251 for configuring the pipeline. A default pipeline can be used or a specific
252 configuration can be indicated by the *Stakeholders* or *Application Developers*.

253 Note, that the *Configuration Module* can be configured both by the *Moni-*
254 *toring Algorithms* themselves and by the *Stakeholders* & *Application Develop-*
255 *ers*. It can thus be regulated how much autonomy the *Learning Engine* has.
256 Moreover, the possibility of human intervention avoids unnecessary learn-
257 ing steps in case the new knowledge, which should be added to the ontology
258 based on the observation from the *Monitoring Algorithms*, is straightforward.
259 Finally, the results of the *Monitoring Algorithms* give the *Stakeholders* & *Ap-*
260 *plication Developers* insight into the behavior and requirements of the users.

261 3.3. *Learning engine*

262 The Pipes-and-Filters architectural design pattern (Bass et al., 2003) was
263 used to design the *Learning Engine*. This data-driven pattern divides a
264 larger processing task into a sequence of smaller, independent processing
265 steps, called filters, that are connected by channels, called pipes. Each filter
266 provides a simple interface, namely it receives messages on the incoming pipe,
267 processes them and provides the results to the outgoing pipe. A filter is thus
268 unaware of its position in the pipeline and which filter precedes and follows
269 it. Because all the filters use similar interfaces they can be combined into
270 different pipelines. Filters can thus easily be added, omitted or rearranged.
271 As a result, the architecture becomes very modular, extensible, re-usable and
272 flexible.

273 3.3.1. *Data collection & input conversion*

274 To be able to use a flexible Pipes-and-Filters architecture, the data ex-
275 changed between the filters needs to be expressed in the same format. A
276 format was developed, which allows expressing both the information which
277 is used as input and the knowledge that is obtained as output, e.g., rules.
278 The format is largely based on the Attribute-Relation File Format (ARFF),
279 which is the text file format used by WEKA (Witten et al., 2011).

280 *The Data Collection Module* is responsible for gathering the necessary
281 input information for the *Learning Engine* from the *Persistence Layer*. The
282 *Input Convertor* converts this data to the data format used by the *Learning*
283 *Pipeline*. The *Data Collection Module* and *Input Convertor* cannot be con-
284 sidered as actual filters for two reasons. First, for any use case scenario they
285 will always appear as the first two steps of the pipeline. Second, the input

286 and output format of these modules is dependent on the source from which
287 the information is collected, e.g., a triple store.

288 3.3.2. *Learning pipeline*

289 The *Pre-Processor* contains several modules to clean up the data. For
290 example, the *Remove Outliers* component removes unrealistic entries from
291 the input data, e.g., impossible sensor values. The *Scale* component centers
292 the input values at zero. This is often beneficial for the learning algorithms
293 of various machine learning techniques. *Feature Selection* can be used to
294 reduce the size of the input data set and thus speed up the data mining.
295 Other examples of pre-processing techniques can easily be integrated into
296 the pipeline as new Filters.

297 The cleaned data is then passed to the *Data Mining* component that pro-
298 vides several techniques to discover trends, e.g., classification rules, decision
299 trees, Bayesian networks or clustering. The results of the *Data Mining* are
300 then processed by the *Post-Processor* to derive the actual information which
301 can be added to the ontology, e.g., rules or thresholds can be derived from a
302 decision tree by the *Rule* or *Threshold Extractor*.

303 The conclusions of the *Post-Processor* are studied further by the *Decision*
304 *Module*. To ensure that the *Knowledge Base* does not become inconsistent
305 when the new knowledge is added, i.e., because it contradicts with already
306 defined knowledge, probabilistic relations are defined between the new and
307 existing knowledge. Moreover, this probability also makes clear to the *Stake-*
308 *holders* that the new knowledge has not been confirmed by rigorous evalua-
309 tion yet. The *Probabilistic Relevance Algorithms* are used to determine the
310 initial probability that should be associated with this new knowledge. For

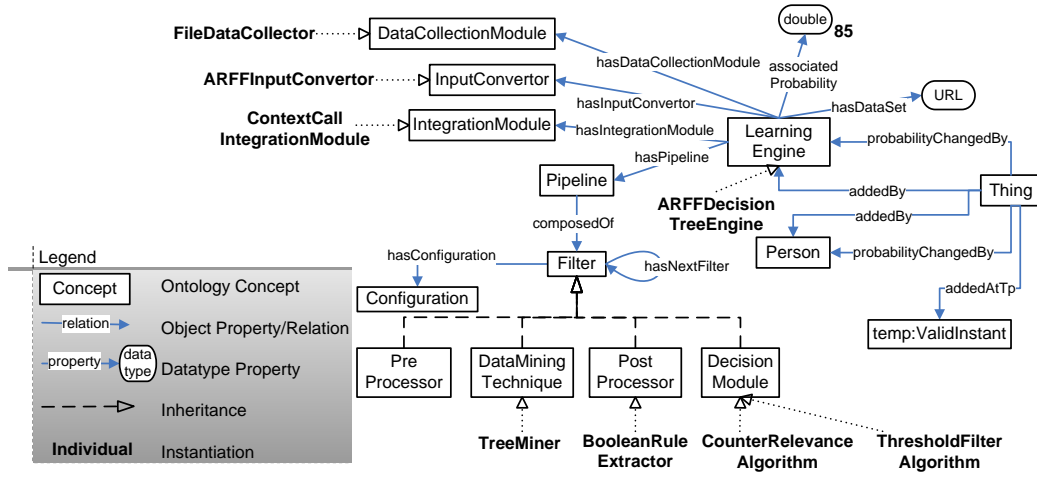


Figure 2: The Learning Ontology

example, it can be calculated how many times a derived rule occurred in the data set on which the data mining technique was trained. However, wrong trends can easily be detected because of skewed or too small data sets. It is also important to only include trends that reflect good and general work practices. Wrong information could clutter the *Knowledge Base* and make the context-aware platform less useable. The *Filter Algorithms* are responsible for detecting and removing these anomalies, e.g., by removing knowledge that received a too low probability by the *Probabilistic Relevance Algorithms*.

The *Learning Pipeline* cannot only be used to learn new information, but also to reassess knowledge that has been previously added to the *Knowledge Base*. In this case, the *Probabilistic Relevance Algorithms* are responsible for in- or decreasing the probability depending on the new information that becomes available about the usage of this knowledge.

3.3.3. Integration module & adapting the probabilities

Finally, the *Integration Module* is responsible for defining the probabilistic relations that connect the new knowledge to the existing knowledge in the *Knowledge Base*. For the same reasons as were already explained in Section 3.3.1 for the *Data Integration* and *Input Convertor Modules*, this module cannot really be considered a filter.

For new knowledge, the probability calculated by the *Probabilistic Relevance Algorithms* is used. When the *Stakeholders* are confronted with a probabilistic decision in their daily work practices, they might be interested in the origin of the information, i.e., how the information was learned, before deciding to follow the recommendation of the context-aware platform or not. Therefore, the *Learning Ontology* was created, which allows associating the learned knowledge with its origin. The most important concepts of this ontology are visualized in Figure 2. This ontology also allows *Application Developers* to easily identify learned knowledge. This enables them to treat this knowledge differently if needed, e.g., ignore it in reliability critical applications or highlight it for the users.

For reassessed knowledge, two thresholds are checked. If the probability calculated by the *Probabilistic Relevance Algorithms* falls below the lowest threshold, the knowledge is removed from the *Knowledge Base* as it is clearly not being used or confirmed by the stakeholders. If the probability exceeds the highest threshold, the knowledge is added to the ontology as generally accepted knowledge, i.e., without an associated probability. Finally, if the probability lies between the two thresholds, the probability of the reassessed knowledge is updated to this probability to reflect its changed reliability. As

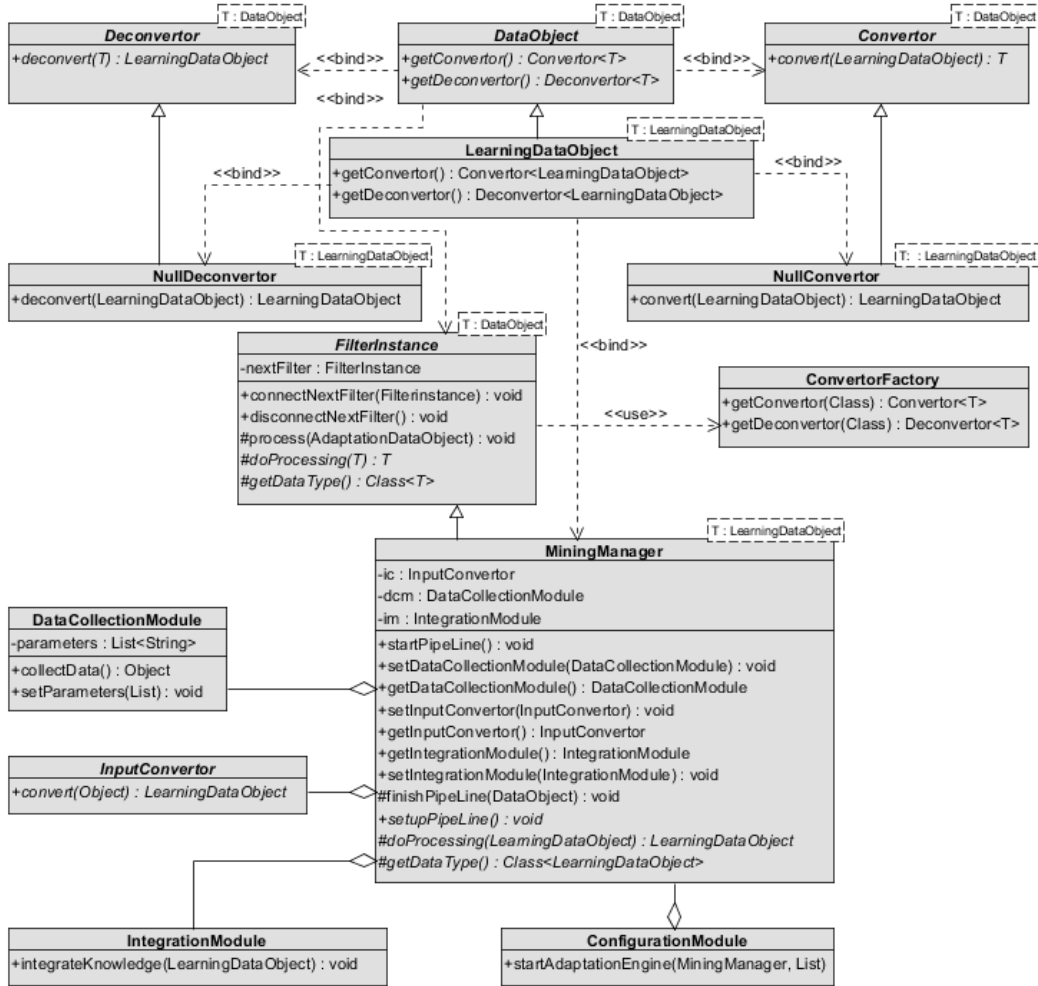


Figure 3: The class diagram of the *Learning Engine* and *Configuration Module*

349 such, a self-learning, context-aware platform is obtained in which knowledge
350 can be added and removed on the fly based on historical information.

351 4. Implementation details

352 The implementation details of the *Context-Aware Platform* are described
353 in Strobbe et al. (2007, 2012b). The platform uses OWL (McGuinness and
354 Harmelen, 2004) as ontology language, Pellet (Sirin et al., 2007) as *DL Rea-*
355 *soner*, Jena Rules (Carroll et al., 2004) and SWRL (Horrocks et al., 2004)
356 to express the *Rules* and SPARQL (Prud’hommeaux and Seaborne, 2008) to
357 query the context information. The platform was extended with the *Proba-*
358 *bilistic Reasoner* Pronto (Klinov, 2008) to enable probabilistic reasoning on
359 the ontologies. Jena is used to manage and persist the ontologies.

360 The *Learning Engine*, *Monitoring Algorithms* and *Configuration Module*
361 were implemented in Java. The class diagram of the *Learning Engine* is
362 visualized in Figure 3. These are the (abstract) classes, which can be used for
363 any scenario. To implement a specific use case, subclasses can be created that
364 implement the specific requirements of the scenario, e.g., a specific pipeline
365 configuration or a specific input convertor. An example of how a specific use
366 case can be implemented is thoroughly explained in Section 5. How these
367 classes can be used to construct and use a specific *Learning Pipeline* with
368 associated *Data Collection Module*, *Input Convertor* and *Integration Module*
369 is visualized with a sequence diagram in Figure 4.

370 As can be seen, the different filters in the *Learning Pipeline* are resp-
371 resented by `FilterInstance` objects. Specific filters, e.g., pre- and post-
372 processors, filter algorithms and data mining techniques, are created as sub-
373 classes of this `FilterInstance` class by implementing the `doProcessing`
374 method. This method specifies how the data is processed by the specific
375 filter, e.g., a `ScalingFilter` that scales the data or a `ClusterFilter` that

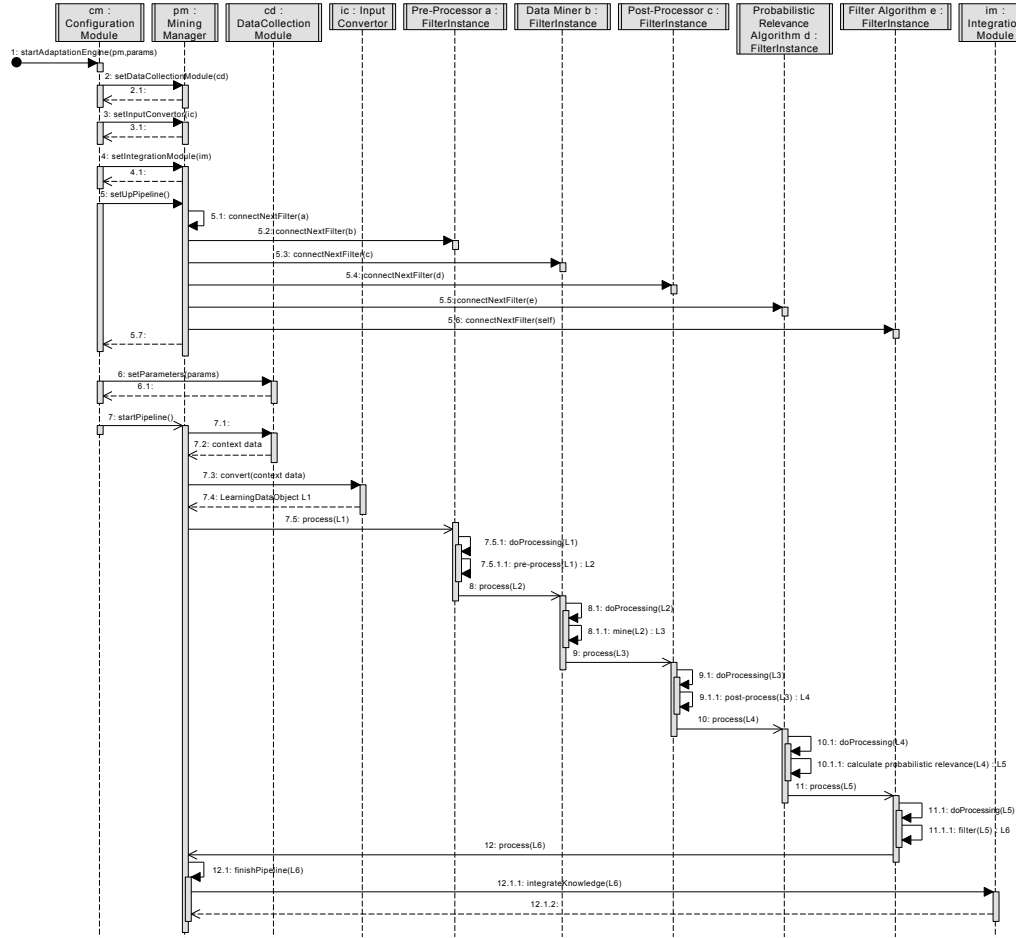


Figure 4: Sequence diagram illustrating the construction and usage of a *Learning Pipeline* with associated *Data Collection Module*, *Input Converter* and *Integration Module*

376 clusters it.

377 As mentioned previously, the data exchanged between the filters in the
 378 pipeline uses the same data format, which is represented by the **Learning-**
 379 **DataObject** Java-Object. This object contains the information about the
 380 different attributes, i.e., ontology concepts, which will be mined, and their

381 data instances. However, to enable logging of the data at any point during
382 the pipeline, this object can easily be serialized to XML.

383 As can be seen, (de)convertors can be used to translate the specific data
384 format to other formats. This is not only necessary to convert the context
385 data gathered by the *Context-Aware Platform* to the data format used by
386 the pipeline, but also to allow the usage of external libraries, e.g., WEKA for
387 data mining. The (de)convertors allow to transform the `LearningDataOb-`
388 `ject` to the format used by the external libraries, e.g., the ARFF format
389 used by WEKA. Each `FilterInstance` indicates which datatype it employs
390 to process the data by using ‘*generic types*’. Based on the indicated type,
391 the framework is able to automatically find the appropriate `Converter` and
392 `Deconverter`. This eases the development of specific use cases and the usage
393 of external libraries. The *Application Developers* only have to develop `Con-`
394 `verter` and `Deconverter` subclasses that implement the conversion to the
395 specific file format used by the `FilterInstance`.

396 To manage the complete pipeline, a special type of `FilterInstance` was
397 created, namely the `MiningManager`. This class is responsible for construct-
398 ing the *Learning Pipeline* out of the separate filters, starting it and processing
399 the results. To implement a specific *Learning Pipeline*, a subclass of the `Mi-`
400 `ningManager` needs to be constructed that implements the `setupPipeline`
401 method. This method initializes the different filters of the pipeline and con-
402 nects them to each other. Each `FilterInstance` is connected to the next
403 `FilterInstance` in the pipeline by using the `connectNextFilter` method.
404 The first `FilterInstance` is connected to the `MiningManager`, while the last
405 `FilterInstance` indicates the `MiningManager` as next filter to ensure proper

406 processing of the result of the *Learning Pipeline*.

407 The `ConfigurationModule` is notified of which data should be collected
408 for the mining process, either by the *Stakeholders* and *Application Developers*
409 or by the *Monitoring Algorithms*. It configures the `MiningManager` to
410 use the appropriate `DataCollectionModule`, `InputConvertor` and `Integra-`
411 `tionModule` that suits this type of data. It also passes the correct parameters
412 to the `DataCollectionModule`, which are needed to retrieve the data from
413 the *Persistency Layer*. Next, the `ConfigurationModule` calls the `setup-`
414 `Pipeline` and `startPipeline` methods of the `MiningManager` to create the
415 pipeline and start the learning process. The latter method first collects the
416 necessary data by using the associated `DataCollectionModule` and converts
417 it to the `LearningDataObject` format with the `InputConvertor`. Next, the
418 `MiningManager` calls the `process` method of the first `FilterInstance` in the
419 pipeline. This `FilterInstance` processes the data with its `doProcessing`
420 method and then calls the `process` method of the next `FilterInstance` in
421 the pipeline. This continues until the last `FilterInstance` calls the `process`
422 method of the `MiningManager`. The `MiningManger` then finishes the learning
423 process by calling the `IntegrationModule` to integrate the knowledge in the
424 *Knowledge Base*.

425 It can be noted that the implemented framework is very extensible, mod-
426 ular and flexible, which allows easy adoption for any use case, as illustrated
427 in the following section.

428 5. Use case: Mining the reasons for patients' call light use to au- 429 tomatically launch calls

430 5.1. Scenario description

431 Nurse call systems are a fundamental technology in continuous care as
432 they are used by caregivers to coordinate work, be alerted of patients' needs,
433 communicate with them through intercoms and request help from other staff
434 members. When patients feel unwell they push a button. The nurses then
435 receive a message with the room number on a beeper. This brings up the
436 question: which nurse goes to the room? The closest one? the one on call,
437 etc.? Current systems often have a very static nature as call buttons have
438 fixed locations, e.g., on the wall next to the bed. There is an increased
439 risk when patients become unwell inside a hallway, staircase or outside as
440 they cannot use the nurse call system. Additionally, the current nurse call
441 algorithms consist of predefined links between beeper numbers and rooms.
442 Consequently, the system presently does not take into account the various
443 factors specific to a given situation, such as the pathology of a patient, e.g.,
444 heart patient or confused, nor the competences of the staff, e.g., nurse or
445 caregiver.

446 The increased introduction of electronic devices in continuous care set-
447 tings facilitated the development of the ontology-based Nurse Call System
448 (oNCS), which allows patients to walk around freely and use wireless nurse
449 call buttons. Additionally, this platform manages the profiles of staff mem-
450 bers and patients in an ontology. A sophisticated nurse call algorithm was
451 developed by the authors. It first determines the priority of the call us-
452 ing probabilistic reasoning algorithms, which take into account the origin

453 of the call and the pathology of the patient. Next, the algorithm finds the
454 most appropriate staff member to handle the call. It dynamically adapts to
455 the situation at hand by taking into account the context, e.g., location of
456 the staff members and patients, the priority of the call and the competence
457 of the different caregivers. The oNCS was implemented according to the
458 *Context-Aware Platform* architecture discussed in Section 3 and visualized
459 in Figure 1. A detailed description of this platform can be found in Ongenae
460 et al. (2011d).

461 The oNCS is also able to automatically launch context calls based on the
462 data generated by the electronic equipment and sensors in the environment,
463 e.g., when a patient spikes a fever or when the light intensity is too high in the
464 room of a patient with a concussion. It is however very difficult for developers
465 to determine in advance all the risky situations for which a context call should
466 be launched. These parameters and their thresholds are very dependent on
467 the specific environment where the oNCS is deployed. Moreover, some of the
468 relations between parameter measurements and calls made by the patient
469 might not even be directly apparent to the caregivers as these relations are
470 not rigorously studied.

471 To detect relations between the parameter measurements and the calls
472 made by patients, the oNCS was extended with *Monitoring Algorithms*, the
473 *Configuration Module* and *Learning Engine*. To evaluate this extension, a
474 relation was simulated and it was investigated whether the *Learning Engine*
475 was able to detect this trend and add it to the *Knowledge Base*. The trend
476 that patients make a call when they exhibit symptoms for Systemic Inflam-
477 matory Response Syndrome (SIRS) (Davies and Hagen, 1997; Nyström, 1998)

478 was chosen as simulated relation. This medically relevant use case could be
479 easily generated, but is challenging for the *Learning Engine* to detect. SIRS
480 is a generalized inflammatory reaction of the organism to a severe medical
481 condition such as acute pancreatitis, severe burn injury, trauma, surgical
482 procedure or infection. If SIRS is the response to an infection, the patient
483 is diagnosed with sepsis. Sepsis has a high mortality rate (30%-40%). The
484 criteria for diagnosing a patient with SIRS are:

- 485 • Tachycardia: heart rate > 90 beats per minute (bpm)
- 486 • Fever or hypothermia: body temperature $> 38^{\circ}\text{C}$ or $< 36^{\circ}\text{C}$
- 487 • Tachypnea: arterial partial pressure of carbon dioxide (PaCO_2) < 32
488 mmHg
- 489 • White Blood Cell (WBC) count $< 4,000$ cells/ mm^3 or $> 12,000$ cells/ mm^3

490 For the diagnosis of SIRS, two or more of these criteria must be fulfilled. This
491 is a challenging scenario for the *Learning Engine* as it involves both param-
492 eters measured at regular intervals by sensors, i.e., the heart rate and body
493 temperature, as well as parameters obtained through the analysis of a blood
494 sample by the laboratory, i.e., WBC and PaCO_2 . Moreover, a combination
495 of conditions needs to be fulfilled before the call should be launched.

496 The following sections illustrate how the *Learning Engine* was imple-
497 mented and the *Learning Pipeline* was constructed, using the (abstract)
498 classes discussed in Section 4, to detect this relation and add it to the *Knowl-*
499 *edge Base*. The resulting pipeline is visualized in Figure 5.

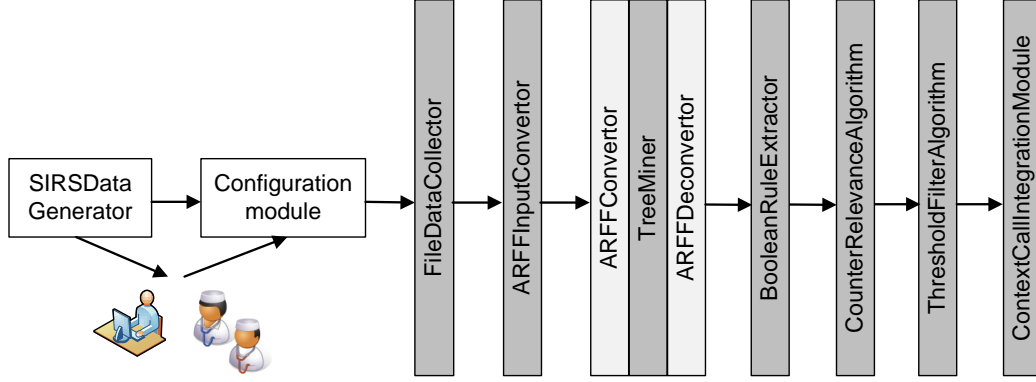


Figure 5: The pipeline used by the *Learning Engine* to tackle the SIRS use case

5.2. Scenario implementation

5.2.1. Generating the SIRS data

To realize the scenario, a dataset needs to be generated in which the trend can be detected that patients make calls when they exhibit SIRS symptoms. This dataset consists of a set of instances, each consisting of five data values, namely a value for the four SIRS parameters and whether or not a call was made. A *SIRS Instance* is defined as an instance, which consists of a combination of the four SIRS parameters that fulfills two or more SIRS criteria. Logically, a *Non-SIRS Instance* is defined as an instance, which fulfills at most one SIRS criterion at a time.

When the different instances are generated, each instance has 15% chance of being a *SIRS Instance*. The parameter values are randomly generated, while ensuring that at least two parameters fulfill the SIRS criteria for *SIRS Instances* and at most one criterion is fulfilled for *Non-SIRS Instances*. The values are generated within realistic bounds, e.g., temperature must be lower

515 than 43 °C. Whether the *SIRS Instance* fulfills two, three or four criteria
516 and whether the *Non-SIRS Instance* fulfills one criterion or none, is also
517 randomly chosen.

518 Finally, each instance needs to be associated with a context call or not.
519 To achieve a realistic dataset, noise is introduced by wrongly classifying the
520 instances, i.e., associating *Non-SIRS Instances* with a call and vice versa. A
521 noise percentage of x means that each *Non-SIRS Instance* has $x\%$ chance of
522 being associated with a call and vice versa.

523 Some example instances are illustrated in Table 1. The first four instances
524 are *Non-SIRS Instances*, while the latter four are *SIRS Instances*. The pa-
525 rameter values that fulfill SIRS criteria are indicated in italic. The calls
526 marked with a *-symbol represent noise. A *Data Generator* was written to
527 create the needed instances and provide them in the ARFF format, i.e., the
528 data format used by WEKA. The resulting file is stored in the *Persistence*
529 *Layer*.

530 5.2.2. The *oNCS* and continuous care ontologies

531 As mentioned in Section 2, little work has been done on the development
532 of high-level ontologies, which can be used to model context information and
533 knowledge utilized across the various continuous care settings, e.g., hospitals,
534 homecare and residential care settings. Therefore we developed the *Continu-*
535 *ous Care Ontology*, which models the generic context information gathered by
536 the various sensors and devices, the different devices, the various staff mem-
537 bers and patients and their profile information, medical conditions, roles and
538 competences and the variety of tasks that need to be performed. A detailed
539 description of this ontology can be found in Ongenae et al. (2011a). The

Heart rate	Body temperature	PaCO ₂	WBC count	Call
61.42	38.62	34.54	4969	No
78.55	37.47	32.68	7746	No
88.37	35.76	46.53	7253	Yes*
67.92	36.10	42.53	12096	Yes*
66.63	40.95	30.56	3740	Yes
91.59	36.78	29.94	12301	No*
94.52	40.67	28.89	4866	Yes
95.23	35.93	31.61	8737	No*

Table 1: Some example instances of the SIRS dataset

540 most important classes of these ontologies pertaining to the use case are vi-
 541 sualized in Figure 6. This ontology references the *Galen Common Reference*
 542 *Model* (Rector et al., 2003) as *Medical Ontology*. The concepts from the *Galen*
 543 *Common Reference Model* are preceded by the **galen** namespace in Figure 6.
 544 The concepts preceded with the **temporal** namespace are imported from the
 545 *SWRLTemporalOntology* (O'Connor and Das, 2010) to represent temporal
 546 information. Finally, the *Wireless Sensor Network (WSN) Ontology* (Ver-
 547 stichel et al., 2010) was imported, as shown by the concepts preceded by
 548 the **wsn** namespace, to represent the knowledge pertaining to observations
 549 made by sensors. The *Probabilistic Domain Ontology* then models the spe-
 550 cific properties of the environment where the oNCS is deployed, e.g., the
 551 specific roles and competences of the staff members and how they map on

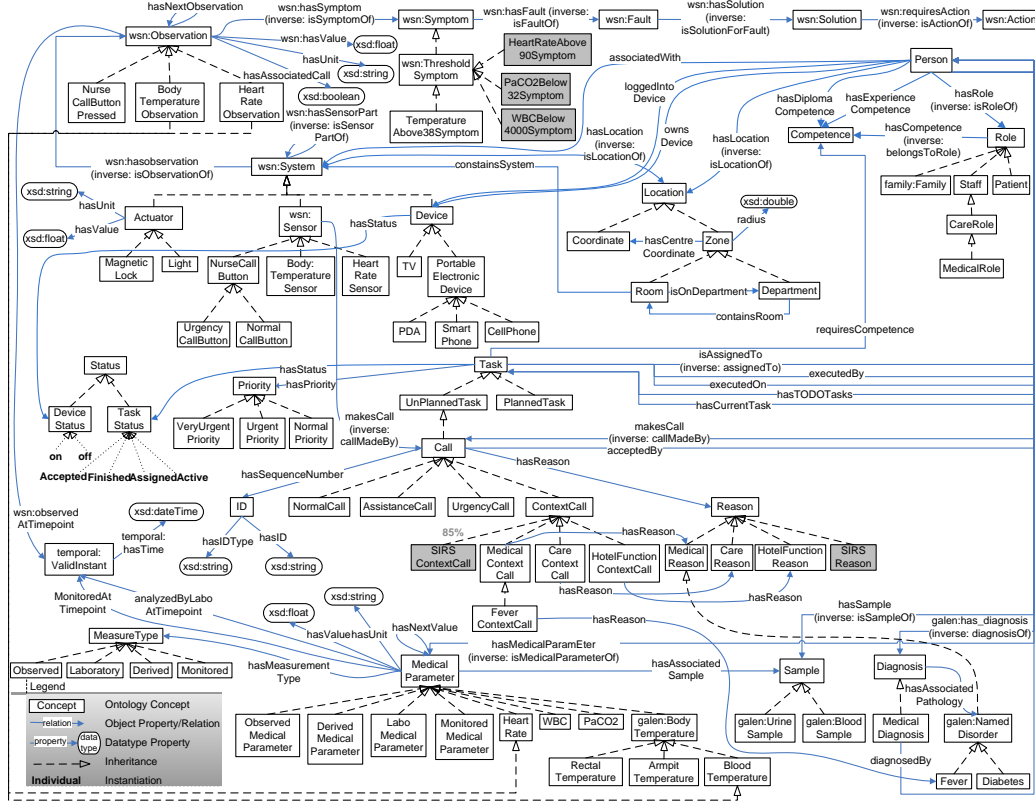


Figure 6: Prevalent classes of the *Continuous Care, Medical and Probabilistic Domain* ontologies pertaining to the SIRS use case

each other.

As can be seen, the model contains a **System** concept which models a system and its components. The ontology allows interpreting the data values monitored by the sensors. For this the ontology uses an observation pattern. A data value monitored by a system is modeled in the ontology as an **Observation**. Rules and axioms added to the ontology allow detecting specific phenomena in these observations, which are modeled as **Symptom** concepts. For example, the **TemperatureAbove38Symptom** class is defined as follows:

BodyTemperatureObservation AND \exists hasValue “ > 38”

560 This axiom ensures that a **BodyTemperatureObservation** of more than
561 38 °C is reclassified as a **TemperatureAbove38Symptom**. Similarly, symptoms
562 can also be reclassified as faults and even as solutions and actions that should
563 be taken.

564 People are modelled through the **Person** concept and their roles and
565 competences can be indicated. It can also be indicated with which person
566 the sensors are associated through the **associatedWith** relationship. The
567 medical parameters collected about a patient, either by sensors, the obser-
568 vations of staff members or the analysis of blood samples, are modelled as
569 **MedicalParameters**. Similar to observations, these parameters can also be
570 reclassified as symptoms. The medical condition of a person can also be
571 modeled, e.g., **Fever**.

572 To model the daily activities of the caregivers and patients, the **Task**
573 concept is used, which is further divided into planned and unplanned tasks.
574 Each task can be assigned a **Status**, e.g., **Active** or **Finished**, a **Priority**
575 and the competences which are needed to execute the task. People can be
576 connected to the tasks through various relationships, e.g., **hasCurrentTask**,
577 **isAssignedTo** or **executedBy**. A **Call** is modelled as an **unPlannedTask**. A
578 call can be associated with a **Reason**, e.g., **Fever**. Four types of calls can be
579 discerned. A **NormalCall** is a call made by a patient, while an **Assistance-**
580 **Call** is launched by a caregiver to request help from another staff member.
581 An **UrgencyCall** is only used for emergency situations, e.g., when a patient
582 needs to be reanimated. Finally, a **ContextCall** is call that is automatically

583 generated by the oNCS as a consequence of certain conditions being fulfilled.

584 Consider for example the following Jena rule:

```
[FeverContextCall:
(?symp rdf:type oncs:TemperatureAbove38Symptom)
noValue(?symp task:hasAssociatedCall)
(?symp wsn:isObservationOf ?system)
(?kind rdf:type oncs:FeverContextCall)
→
createContextCall(?system, ?kind)
(?symp task:hasAssociatedCall 'true'~xsd:boolean)]
```

585 The first line represents the name of the rule. First, it is sought if a
586 body temperature of more than 38 °C was observed for which a call has not
587 been launched yet. Next, the system that made the observation is retrieved.
588 Finally, the type of call that should be created is specified. As a result, the
589 functor `createContextCall` is called, which creates a `ContextCall` of type
590 `FeverContextCall` and associates the system that made the observation with
591 this call. The functor also assigns the status `Active` to the call. Moreover,
592 the `hasAssociatedCall` relationship is set to `true` to make sure that the
593 rule does not fire again.

594 The oNCS contains rules that fire when active calls are added to the
595 ontology. Based on the context information, these rules assign the most ap-
596 propriate staff member to the call. More information about these assignment

597 rules can be found in Ongenae et al. (2011d).

598 Similar to how the fever example was modeled, the SIRS use case can
599 be easily represented using these classes. Individuals of type **BodyTempera-**
600 **tureSensor** and **HeartRateSensor** are created to represent the sensors that
601 measure the medical parameters of the patients. These sensors make obser-
602 vations of type **BodyTemperatureObservation** and **HeartRateObservation**
603 respectively, which are associated with their sensors through the **hasObser-**
604 **vation** relation. The measured value is indicated with the **hasValue** relation.
605 Individuals of type **BloodSample** are created, that represent the blood sam-
606 ples analyzed by the laboratory to determine the WBC count and PaCO₂ of
607 the patient. These results are captured in the ontology as medical parameters
608 of type **WBC** and **PaCO2**. They are associated with their blood sample through
609 the **hasAssociatedSample** relationship. Finally, when a patient makes a call
610 by pressing a button, an individual of type **Call** is created in the ontology,
611 which is connected through the **callMadeBy** relationship with the patient.
612 Through reasoning, this call is reclassified as a **NormalCall** as it is made by
613 a person with as role **Patient**.

614 A mobile nurse call application was also developed, which is used by the
615 caregivers to receive, assess and accept, i.e., indicate that they are going to
616 handle, calls. A nurse can also use the application to contact other staff
617 members or the patient, e.g., to request the reason for the call or to give
618 feedback. Before a nurse is able to indicate a call as finished, the reason for
619 the call must be indicated either on the mobile application or the terminal
620 next to the bed of the patient. This reason is also entered in the ontology.
621 The mobile application is further explained in Ongenae et al. (2011c).

622 5.2.3. Collecting the data & input conversion

623 As the data is generated, no *Monitoring Algorithms* are needed. However,
624 a *Monitoring Algorithm* could easily be written as follows. Relationships
625 need to be found between medical parameters of patients and the calls that
626 they make. The *Context Call Monitoring Algorithm*, monitors the ontology
627 for calls of type **NormalCall**. When such a call is added to the ontology,
628 the algorithm collects the most recent value for each medical parameter that
629 is measured about the patient who made the call. This information can
630 easily be retrieved using SPARQL queries. As not every medical parameter
631 is measured for every patient, the dataset possibly contains missing values.
632 When the call has been completely handled by the caregiver, the algorithm
633 also retrieves the reason, which was attached to the call. As such, different
634 data sets can be created, grouping calls together which have similar reasons.
635 These datasets can differ in granularity of the reason. For example, a dataset
636 could be created for all the calls with a **MedicalReason** or for all the calls
637 with the more specific reason **Fever**. All calls of the second dataset would
638 also be part of the first dataset, as **Fever** is a subclass of **MedicalReason**.
639 Each of these datasets could be used as input for the *Learning Engine*. Other
640 ways of grouping the data instances can also be employed, e.g., grouped per
641 patient or grouping the instances of patients that have a similar pathology.
642 The *Context Call Monitoring Algorithm* keeps track of how many instances
643 have been collected for each dataset. When a representative amount has been
644 gathered, the dataset is expanded with negative examples. For example, the
645 medical parameters of the patients already present in the dataset can be
646 collected at a timepoint when they have not seen a caregiver or made a call

647 for a while or at a timepoint they made a call for a different reason. Finally,
648 the *Monitoring Algorithms* invoke the *Configuration Module* to initiate the
649 *Learning Engine*. The datasets can also be intermediately shown to the
650 *Stakeholders* and *Application Developers* for inspection. In this use case, the
651 *Data Generator* takes on the role of the *Monitoring Algorithm*.

652 The *Monitoring Algorithms* can store the datasets in the *Persistence*
653 *Layer* in a format that best suits their needs. For the *Data Generator*,
654 the ARFF format was chosen. Ontology individuals could also be directly
655 stored in a triple store. The *Monitoring Algorithm* or the *Data Generator*
656 indicates the location of the data and its format to the *Configuration Mod-*
657 *ule*. They also indicate which `MiningManager` should be used to process the
658 data. Different types of *Learning Pipelines*, which each consist of a combina-
659 tion of filters that suit the needs of a particular use case, can be created by
660 implementing several subclasses of the `MiningManager`. This allows multiple
661 *Monitoring Algorithms* to run at the same time and the collected data to be
662 processed by the `MiningManager`, and thus *Learning Pipeline*, that matches
663 with the goal of the algorithm.

664 The *Configuration Manager* configures the `MiningManager` to use the ap-
665 propriate `DataCollectionModule` and `InputConvertor` that suits the format
666 of the data. The subclass `FileDataCollector` of the `DataCollectionMod-`
667 `ule` class was implemented, which is able to read the data from a file at
668 a specified location. The result is a `String`, which is provided to the ap-
669 propriate `ARFFInputConvertor`. This subclass of `InputConvertor` is able
670 to translate this `ARFF-String` to the `LearningDataObject` format, which is
671 used by the *Learning Pipeline*. During the translation it also checks if the

specified value for an attribute, e.g., 38 for the body temperature parameter, is compatible with the type of this attribute. For example, it is not allowed to assign a **String** to a numerical attribute. Illegal data instances are discarded.

5.2.4. Mining the sensor data using a C4.5 decision tree

A *Pre-Processing* filter was not implemented for this use case, as it works on generated data. If the previously discussed *Context Call Monitoring Algorithm* was used, several *Pre-Processing* filters could be used. For example, a **RemoveOutliers** filter could be employed to remove outliers or impossible parameter values in the dataset. Moreover, the number of features, i.e., measured medical parameters, in the dataset would be relatively high. A **FeatureSelection** filter could be used to select the most interesting features for the *Data Mining* step. Finally, a **MissingValues** filter would be able to deal with the missing values in the dataset.

The *Data Mining* filter needs to find relations in the generated dataset between the sensor measurements and the occurrence of a call. Supervised (Witten et al., 2011) classification techniques (Kotsiantis, 2007) consider a set of input attributes, e.g., the different sensor types, and a output attribute, also called the class attribute or the label, e.g., whether a call was made or not. These techniques then try to build a model that fits this data set and derives relationships between the input attributes and the label. Building a decision tree (Kotsiantis, 2013) based on the information captured in the dataset is a well-known and easy supervised classification technique. A decision tree is a tree structure in which each leaf represents a value that the label can assume, e.g., Yes or No. The internal nodes of the tree represent the attributes

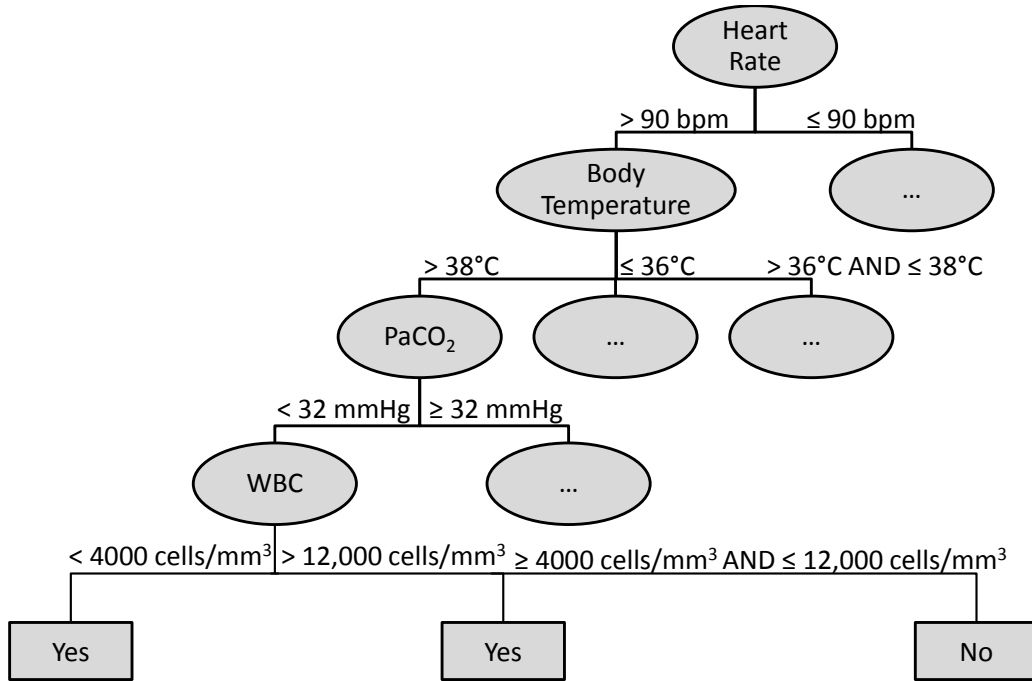


Figure 7: Part of the decision tree of the SIRS example

on which a decision is based, while the branches represent conditions that the attributes need to fulfill. As an example, a part of the decision tree of the SIRS example is shown in Figure 7. To determine the label of a certain data instance, one just needs to follow the tree from the root to the leaves along the branches for which the instance fulfills the conditions. Essentially, a decision tree forms a compact representation of classification rules. For example, the decision tree shown in Figure 7 contains the classification rule:

$$\text{HeartRate} > 90 \text{ bpm AND BodyTemperature} > 38^{\circ}\text{C AND} \\ < 32 \text{ mmHg AND} < 4000 \text{ cells/mm}^3 \rightarrow \text{Yes}$$

Different techniques can be used to build such a decision tree out of a data

705 set, e.g., the Iterative Dichotomiser 3 (ID3) (Quinlan, 1986) or C4.5 (Quinlan,
706 1993) algorithm. The latter is a more sophisticated algorithm as it allows
707 that attributes have numeric values (Quinlan, 1996), is able to handle missing
708 values and prunes the tree in order to make it more compact and avoid
709 overfitting (Everitt and Skron dal, 2010).

710 To implement the C4.5 decision tree, an external library is used, namely
711 WEKA. WEKA provides its own implementation of the C4.5 algorithm,
712 namely J4.8, which was used in this research. A subclass of the `FilterInstance`
713 abstract class was implemented, called `TreeMiner`. As pre-
714 viously mentioned, WEKA uses the ARFF data format to represent data.
715 Therefore, an `ARFFDataObject` was created as a subclass of `DataObject` and
716 (de)convertors were implemented that are able to translate the internal data
717 format of the *Learning Engine*, i.e., `LearningDataObject`, to and from the
718 ARFF data format. As mentioned in Section 4, it is enough to indicate in
719 the `TreeMiner` that the filter uses the `ARFFDataObject` in the `getDataType`
720 method and the framework will automatically use the correct (de)convertors
721 to transform the data. Which attribute should be used as label can be indi-
722 cated in the `TreeMiner` class. In case the label is not indicated, the `TreeM-`
723 `iner` assumes that the last attribute in the data format is the label. The
724 `ARFFInputConvertor`, discussed in the previous section, makes sure that the
725 last attribute is indeed the label. The `doProcessing` method then calls the
726 Java API of WEKA to build the decision tree. However, the J4.8 algorithm
727 does not allow to retrieve separates branches and nodes of the tree. Only
728 a textual representation of the complete decision tree can be obtained. For
729 example, the textual representation of the tree visualized in Figure 7 is:

```

N0 [label="HeartRate" ]
N0 → N1 [label=" > 90"]
N1 [label="BodyTemperature" ]
N1 → N2 [label=" > 38"]
N2 [label="PaCO2" ]
N2 → N3 [label=" < 32"]
N3 [label="WBC" ]
N3 → N4 [label=" < 4000"]
N4 [label="Yes"]

```

730 The nodes and branches are identified and translated to the **Learning-**
731 **DataObject** format such that the result can be forwarded to the next step
732 in the pipeline. It is important to note that new results are always added to
733 the data being exchanged, so that the original data set also stays available
734 for the following steps in the pipeline.

735 The *Post-Processing* filter is responsible for deriving the rules out of the
736 textual representation of the decision tree provided by the J4.8 algorithm.
737 Therefore, the **BooleanRuleExtractor** subclass of the **FilterInstance** class
738 was implemented. The implemented **doProcessing** method takes into ac-
739 count that the label has a boolean value, i.e., Yes or No. Only the branches
740 that result in a positive leaf need to be translated into a rule, as only those
741 rules will result in calls. The **doProcessing** method starts from a positive
742 leaf and follows the branches until the root is reached. Each branch that is
743 crossed is added as a condition in the rule. The iterative build-up of the rule

744 according to the output of the J4.8 algorithm illustrated in Figure 7 is as
745 follows:

Step 1: $\rightarrow Yes$

Step 2: $WBC < 4000 \rightarrow Yes$

Step 3: $PaCO_2 < 32 \text{ AND } WBC < 4000 \rightarrow Yes$

Step 4: $BodyTemperature > 38 \text{ AND } PaCO_2 < 32 \text{ AND}$
 $WBC < 4000 \rightarrow Yes$

Step 5: $HeartRate > 90 \text{ AND } BodyTemperature > 38 \text{ AND}$
 $PaCO_2 < 32 \text{ AND } WBC < 4000 \rightarrow Yes$

746 The resulting rules are represented in the `LearningDataObject` format
747 such that they can be processed by the *Decision Module*.

748 5.2.5. Filtering and integrating the rules

749 As mentioned in Section 3, probabilities are attached to the discovered
750 rules to express their reliability to the users and to ensure that the *Knowledge*
751 *Base* remains consistent, i.e., that the new knowledge does not contradict
752 already existing knowledge.

753 To calculate the initial probability, the `CounterRelevanceAlgorithm` was
754 implemented as a subclass of the `FilterInstance` class. This algorithm ap-
755 plies the rule to the original dataset, which is still included in the `Learning-`
756 `DataObject`. The percentage of times that the rule labels the data correctly,
757 i.e., the conditions of the rule are fulfilled and the label is Yes, is used as
758 probabilistic value. As the data for this use case was generated, this prob-
759 ability thus reflects the amount of noise in the dataset. For the remainder

760 of the text, it assumed that the rule, which was presented in the previous
761 section, receives a probability of 85%.

762 A simple filter algorithm, namely the `ThresholdFilterAlgorithm` was
763 implemented as subclass of the `FilterInstance` class. This algorithm filters
764 the rules for which the probability is lower than a specified probability, e.g.,
765 50%. This rule is thus not added to the *Knowledge Base*. However, the rule
766 and its associated probability is archived in the *Persistence Layer*.

767 Finally, the `ContextCallIntegrationModule`, a subclass of the `Inte-`
768 `grationModule` class, is responsible for integrating the rules and associated
769 probabilities in the *Knowledge Base*. First, new subclasses of `ContextCall`
770 and `Reason` are introduced in the ontology, with as name the condition of
771 the rule added before the suffix `ContextCall` and `Reason` respectively. For
772 brevity, `SIRSContextCall` and `SIRSReason` are used to refer to the concepts
773 that are created for the rule, which is used as running example, i.e., the rule
774 that fulfills each of the four criteria. These concepts are visualized in grey in
775 Figure 6. Pronto is used to represent and reason on the probabilistic infor-
776 mation in the ontology. To express generic probabilistic knowledge, Pronto
777 uses Generic Conditional Constraints (GCCs) (Lukasiewicz, 2007). Generic
778 means that the knowledge does not apply to any specific individual but rather
779 to a fresh, randomly chosen one. A GCC is of the form $(D \text{---} C)[l, u]$ where D
780 and C are classes in the ontology and $[l, u]$ is a closed subinterval of $[0, 1]$. To
781 represent these GCCs in the ontology, Pronto employs subsumption axiom
782 annotations. For example, to express the fact that the `SIRSContextCall` is
783 a `ContextCall` with only 85% probability, the following subsumption axiom
784 annotation is added to the ontology:


```

< owl11:Axiom >
< rdf:subject rdf:resource="#SIRSContextCall" >
< rdf:predicate rdf:resource="&rdfs;subClassOf" >
< rdf:object rdf:resource="#ContextCall" >
< pronto:certainty > 0.85;1 < /pronto:certainty >
< owl11:Axiom >

```

785 Second, a **Symptom** concept is created for each parameter condition in the
786 discovered rule, for example **HeartRateAbove90Symptom**, **BodyTemperature-**
787 **Above38Symptom**, **PaCO2Below32Symptom** and **WBCBelow4000Symptom**. These
788 classes are defined by axioms, for example the **HeartRateAbove90Symptom** is
789 defined as:

HeartRateObservation AND \exists hasValue “ > 90”

790 If a class with a similar definition already exists, the existing class is used.
791 This can be checked by searching for equivalent classes in the ontology with
792 a Reasoner. For example, **BodyTemperatureAbove38Symptom** is not added
793 to the ontology, as **TemperatureAbove38Symptom** is an equivalent class. The
794 newly created **Symptom** classes are visualized in grey in Figure 6.

795 Third, the rules are translated to a Jena Rule using the created classes
796 and added to the *Knowledge Base*. For example, the rule from the previous
797 section is translated to four Jena Rules. For example, the following Jena
798 Rule launches when all the requirements are met and at least one of the
799 symptoms does not have an associated call yet:

```

[SIRSContextCall:
(?symp1 rdf:type oncs:HeartRateAbove90Symptom)
noValue(?symp1 wsn:hasNextObservation)
(?symp2 rdf:type oncs:TemperatureAbove38Symptom)
noValue(?symp2 wsn:hasNextObservation)
(?symp3 rdf:type oncs:PaCO2Below32Symptom)
noValue(?symp3 medical:hasNextValue)
(?symp4 rdf:type oncs:WBCBelow4000Symptom)
noValue(?symp4 medical:hasNextValue)
noValue(?symp1 task:hasAssociatedCall)
(?symp1 wsn:isObservationOf ?system)
(?kind rdf:type oncs:SIRSContextCall)
→
createContextCall(?system, ?kind)
(?symp1 task:hasAssociatedCall 'true'^^xsd:boolean)]
(?symp2 task:hasAssociatedCall 'true'^^xsd:boolean)]
(?symp3 task:hasAssociatedCall 'true'^^xsd:boolean)]
(?symp4 task:hasAssociatedCall 'true'^^xsd:boolean)]

```

800 For each symptom a rule is created. The only difference between the
801 rules is that the condition for an associated call is checked for a different

802 symptom each time. This is because the different symptoms on their own
 803 might already have launched context calls for other reasons, e.g., the **Temper-**
 804 **atureAbove38Symptom** might already have launched a **FeverContextCall**.
 805 Afterwards, all the symptoms are associated with a call to ensure that only
 806 one context call is launched. The rule also ensures that the most recent
 807 parameter values are taken into account by checking whether there are no
 808 next observations or parameter values through the **hasNextObservation** and
 809 **hasNextValue** relations.

810 When the rule is fulfilled, a new context call is added to the *Knowledge*
 811 *Base*. Consequently, the oNCS will detect the new context call and assign
 812 a staff member to it. The Pronto reasoner can then be used to retrieve the
 813 probabilistic information associated with the call. This information can then
 814 be conveyed to the assigned caregiver through the mobile application.

815 As only subclasses are added to the ontology and no knowledge is re-
 816 moved, it is unlikely that the ontology will become inconsistent. However,
 817 if the ontology does become inconsistent, the following solution can be em-
 818 ployed. When new information is added to the ontology, the consistency is
 819 checked. If the ontology is no longer consistent, the information is identified
 820 with which the new knowledge conflicts. Pronto allows that different chunks
 821 of probabilistic information conflict with each other. For example, a bird is
 822 flying object with high probability and all penguins are birds, but a penguin
 823 has a low probability of flying. More specific probabilistic constraints are thus
 824 allowed to override more generic ones. The conflicting information is anno-
 825 tated with the probabilistic interval $[1,1]$, which indicates that the knowledge
 826 is generally true. Consequently, we are now dealing with conflicting, proba-

827 bilistic knowledge and the rule of increasing specificity can be employed to
828 resolve the conflict. As such, we ensure that the ontology remains consistent.

829 Finally, the *Integration Module* also associates the learned knowledge with
830 information about the *Learning Engine* that created it by using concepts
831 from the *Learning Ontology*. The individuals, which are created to realize
832 this goal, are visualized in bold in Figure 2.

833 Note that **ContextCall**, **Symptom** and **Reason** concepts and an associated
834 probabilistic annotation axiom and Jena Rule are created for each discovered
835 rule.

836 5.2.6. *Adapting the probabilities*

837 This step was not implemented as it requires the system to be deployed
838 such that information about the usage of the new knowledge by the caregivers
839 can be acquired. However, it is briefly discussed how this task of adapting
840 the probabilities could be realized for this use case.

841 A *Monitoring Algorithm* could be implemented, which takes as parame-
842 ter the newly created context call, e.g., in this case **SIRSContextCall**. The
843 algorithm monitors the *Knowledge Base* and collects calls of this type, which
844 have been launched by the system. For each call, its reason and the symp-
845 toms that caused the calls to be launched are retrieved. When nurses handle
846 calls, they need to input the reason for the call. For context calls, they can
847 affirm the reason, which was assigned by the framework, e.g., SIRS. They can
848 also choose to change it, e.g., to false because the call was unnecessary. As
849 such a dataset is created for each rule, which maps the values of the medical
850 parameters on the associated reason. When a representative amount of data
851 has been collected, this dataset can be retrieved by the **FileDataCollector**

and converted by the `ARFFInputConverter`. The output can then be processed by a *Learning Pipeline* consisting of only one filter. This filter is a *Probabilistic Relevance Algorithm*, which simply calculates the percentage of calls for each rule for which the reason was not changed. This means that caregivers deemed the reason to be correct. This percentage is then given to the *Integration Module*, which adapts the probability for this rule in the ontology to this calculated percentage. As explained in Section 3.3.3, if the calculated percentage exceeds or falls below the probability thresholds specified in the *Integration Module*, the knowledge is removed from the ontology or added as generally accepted knowledge without a probability.

5.3. Evaluation set-up

To evaluate the applicability of the framework, it is important to assess the correctness of the derived rules. The correctness of the used data mining techniques is influenced by the size of the dataset and the amount of noise. To assess the influence of the latter, the *Learning Pipeline* was consecutively applied to datasets of the same size, but with an increasing amount of noise. The amount of noise is varied from 0% to 50% in steps of 1%. As mentioned in Section 5.2.1, a noise percentage of x means that each *Non-SIRS Instance* has $x\%$ chance of being associated with a call and vice versa. It is unnecessary to increase the noise percentage beyond 50% as a random label is assigned at this point and the dataset becomes meaningless. The amount of noise needs to be increased in a dataset of realistic size. The WBC and PaCO₂ parameters are derived by the laboratory by analyzing a blood sample. Consequently, it is unlikely that more than two different values for these parameters will be generated per patient per day. If we assume that a department contains

877 on average 30 patients and that we want to wait at most 28 days before we
878 run the self-learning framework for the first time, a realistic dataset contains
879 1,680 instances, i.e., 30 patients x 28 days x 2 entries per patient per day.

880 The influence of the size of the dataset on the correctness is evaluated by
881 consecutively applying the *Learning Pipeline* to datasets of increasing size.
882 The dataset sizes range from 100 to 2,000 instances in steps of 100 instances.
883 It can be noted that this range also contains the size of the dataset used for
884 the correctness tests that evaluate the influence of noise, i.e., 1,680 instances.

885 It is also important to evaluate the performance, i.e., execution time and
886 memory usage, of the developed *Learning Engine*. Although, the learning
887 process will mostly run in the background, it is important to assess the
888 amount of resource usage. Most healthcare environments have a limited
889 amount of resources and delegating the processing to the cloud is often dif-
890 ficult because of privacy issues. To evaluate the influence of noise on the
891 performance, the same datasets were used as for the correctness tests. How-
892 ever, to assess the influence of the size of the dataset, datasets were generated
893 with sizes ranging from 1,000 to 30,000 in steps of 1,000 instances. Bigger
894 datasets were used as it is important to explore the limits of the proposed
895 framework.

896 To achieve reliable results, each test was repeated 35 times, of which the
897 first three and the last two were omitted during processing. For each run,
898 a new dataset was generated. Finally, the averages across the 30 remaining
899 runs are calculated and visualized in the form of graphs. The tests were
900 performed on a computer with the following specifications: 4096 megabyte
901 (MB) (2 x 2048 MB) 1067 megahertz (MHz) Double Data Rate Type Three

902 Synchronous Dynamic Random Access Memory (DDR3 SDRAM) and an
903 Intel Core i5-430 Central Processing Unit (CPU) (2 cores, 4 threads, 2.26
904 gigahertz (GHz), 3 MB cache).

905 The term detection rate is introduced to assess the correctness. The SIRS
906 use case is detected when the criteria for each of the four parameters of the
907 SIRS use case are discovered. If one or more criteria is not learned, the SIRS
908 use case is considered undetected. The detection rate of a dataset with a
909 particular size is defined as the percentage of the 30 test runs for this size for
910 which the SIRS use case was completely detected. For example, a detection
911 rate of 50% for a dataset of 100 instances means that for 15 test runs of this
912 dataset size the SIRS criteria were detected.

913 To assess the correctness, the relative error of the SIRS criteria is calcu-
914 lated. The relative error expresses how much the learned criterion deviates
915 from the actual SIRS criterion. For example, a relative error of 5% for the
916 “heart rate > 90” criterion indicates that the discovered threshold deviates
917 from 90 by 5%. Note that the body temperature and WBC parameters have
918 both an upper and lower threshold, while the heart rate and PaCO₂ have
919 only one threshold.

920 5.4. Results

921 5.4.1. Correctness

922 Figure 8 depicts the detection rate as a function of the size of the dataset.
923 The detection rate is relatively low for small datasets, but it quickly increases
924 and reaches 100% for a dataset with 800 instances. When a dataset contains
925 at least 1,000 instances, the detection rate is always 100%.

926 The detection rate is off course related to the relative error. In Figure 9




Figure 8: The detection rate of the SIRS use case as a function of the size of the dataset

the relative error is depicted for each of the SIRS criteria as a function of the size of the dataset. A missing value, i.e., the criterion was not learned, corresponds to a relative error of 100%. Consequently, a low detection rate corresponds to high relative error. When the dataset reaches a 1,000 instances and a detection rate of 100% is thus achieved, the relative error stays below 1%. This means that for a dataset of 1,000 instances, the threshold is discovered for each criterion and it only deviates from the actual threshold by at most 1%, which is a very good result. If we consider that the parameters are collected twice a day for each patient in a department with 30 patients, enough instances would be collected after 17 days.

Figure 10 visualizes the relative errors for each of the SIRS criteria as a function of the amount of noise in a realistically sized dataset of 1,680

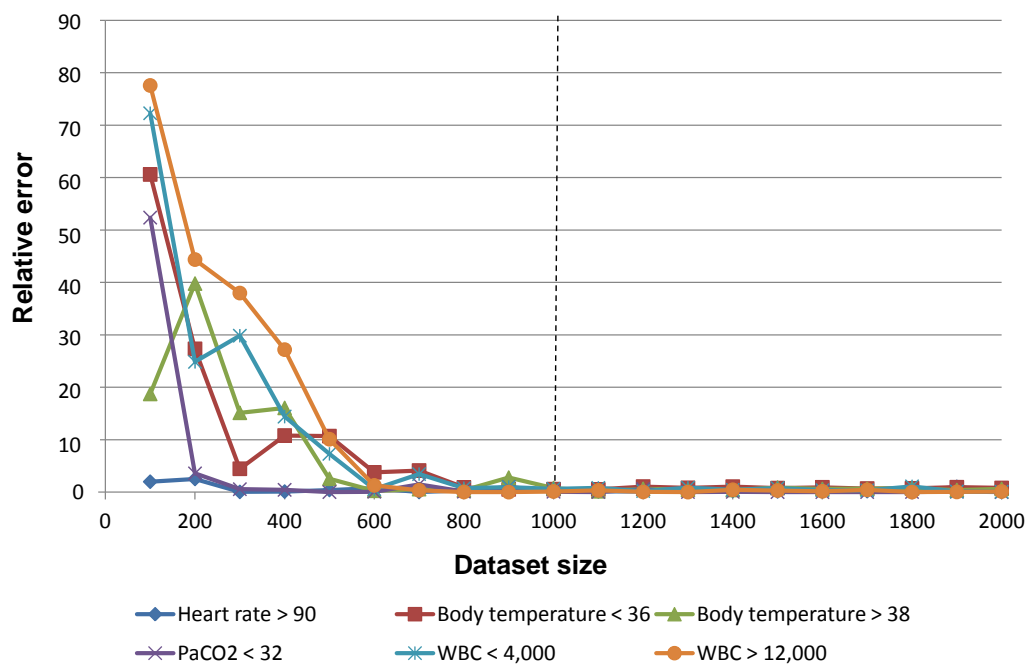


Figure 9: The relative errors for each of the SIRS criteria as a function of the size of the dataset

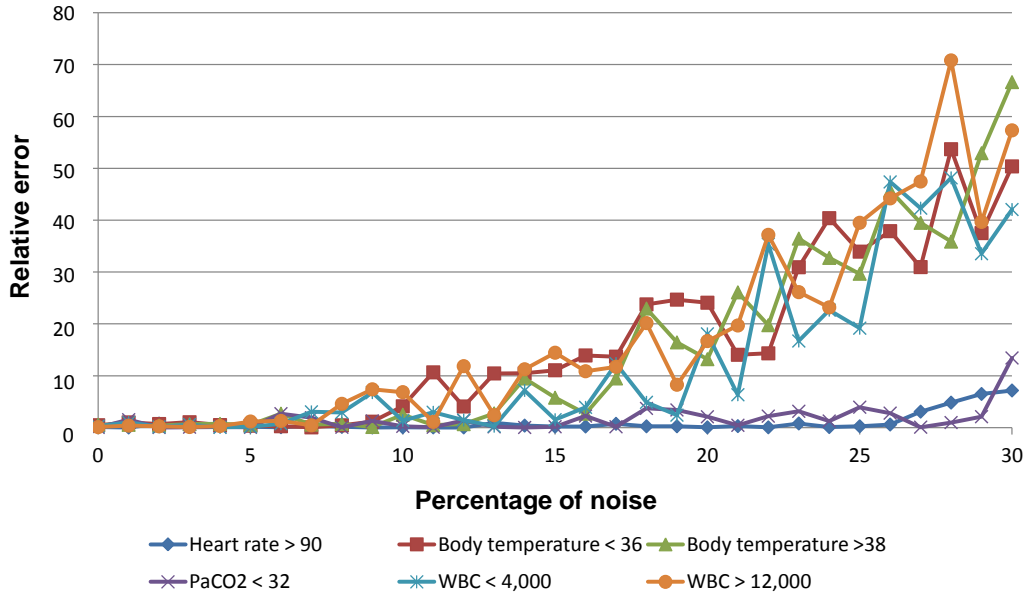


Figure 10: The relative errors for each of the SIRS criteria as a function of the amount of noise in the dataset

instances. It is clear that the *Learning Engine* is insensitive to a noise rate of less than 5%. If the amount of noise increases, the relative errors quickly rise to 10% and higher. A relative error of 10% on the lower threshold of the body temperature, already implies a difference of 3.6 °C. This is unacceptable. In contrast, a relative error of 10% on the lower bound of the WBC only indicates a difference of 400 cells/mm³. The acceptability of the relative error thus depends on the kind and range of the parameter.

5.4.2. Performance

The execution time as a function of the size of the dataset is depicted in Figure 11. Only the execution time of the most relevant pipeline steps




Figure 11: Execution time as a function of the dataset size

949 is shown. The execution time of the CounterRelevanceAlgorithm and
 950 ThresholdFilterAlgorithm is negligible compared to the execution times

951 of the visualized modules. The execution time of the `ContextCallInte-`
 952 `grationModule` depends heavily on the complexity and the amount of data
 953 in the ontology. As the ontology was not initialized with a realistic data
 954 set, e.g., representing a realistic amount of staff members and patients, the
 955 execution time of this module is not shown. The size of the decision tree
 956 build by WEKA depends on the number of attributes in the dataset, but
 957 is independent of the number of instances, i.e., the size of the dataset. As
 958 the number of attributes, namely the four SIRS parameters and the label,
 959 stays constant and the Post-Processing step only processes the model build
 960 by WEKA, the execution time of this step is not influenced by the size of
 961 the dataset. Moreover, the execution time of the `BooleanRuleExtractor`
 962 was also negligible compared to the execution times of the depicted modules.
 963 The processing of the data by WEKA can be split up into two steps, namely
 964 transforming the ARFF format to Java Objects and the actual execution of
 965 the J4.8 algorithm to build the model. The execution times of both these
 966 steps are visualized.


967 It can be derived from Figure 11a that the execution time of the self-
 968 learning framework is linear as a function of the size of the dataset. The
 969 execution of the J4.8 algorithm by WEKA consumes the largest amount of
 970 execution time. It can also be noted that the `ARFFInputConvertor` con-
 971 sumes a considerable amount of execution time. This *InputConvertor* needs
 972 to translate a `String`-based representation of an ARFF-file to the internal
 973 data format used by the *Learning Pipeline*, namely `LearningDataObject`.
 974 Moreover, it needs to check if each value also fulfills the type requirements
 975 of the attribute, e.g., that a `String` is not provided where a numerical value

976 is expected. The `ARFFConverter` and `ARFFDeconverter`, which are used by
977 the Data Mining step to translate the internal data format to and from the
978 ARFF format used by WEKA, are more performant. This is because these
979 converters translate to and from a Java Object representation of the internal
980 format, which is more structured and is thus processed more easily.

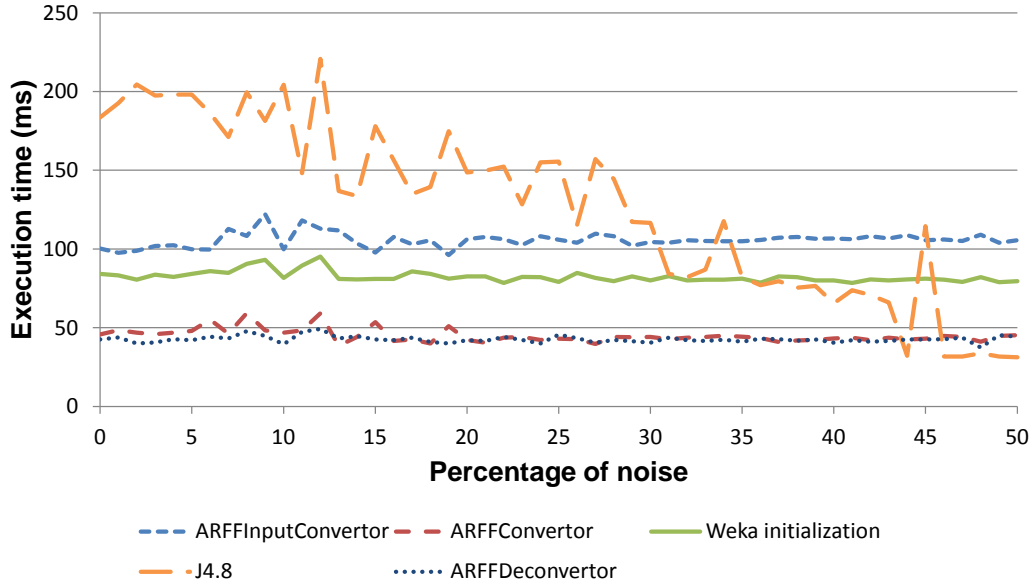
981 Figure 11b illustrates that the execution time of each of the visualized
982 modules is also linear as a function of the size of the dataset. The complexity
983 of the J4.8 algorithm is $O(m * n^2)$ for a dataset with m instances and n
984 attributes (Su and Zhang, 2006). Since the number of attributes is constant
985 in this use case, this reduces to a complexity, which is linear in the number
986 of instances, i.e., $O(m)$. The `ARFFInputConverter`, `ARFFConverter` and
987 `ARFFDeconverter` are also linear in the size of the dataset, as they need to
988 (de)convert all the instances in the dataset one by one.

989 The execution time needed to process the dataset of realistic size, i.e.,
990 1,680 instances, is lower than 100 ms, which is a negligible delay. This means
991 that the monthly patient data of a department with on average 30 patients
992 can be processed very efficiently.

993 Figure 12a depicts the execution time as a function of the amount of
994 noise for the realistic dataset containing 1,680 instances. As the measured
995 execution times are quite small, i.e., lower than 30 ms, the graphs are quite
996 erratic and unpredictable. To get a clear view on the underlying trends, the
997 performance tests were repeated for a dataset consisting of 5,000 instances.
998 The amount of noise in this dataset is also gradually increased. The resulting
999 graph is visualized in Figure 12b. It can be noted that only the execution time
1000 of the J4.8 algorithm is influenced by the amount of noise in the dataset. The



(a) Dataset of 1,680 instances



(b) Dataset of 5,000 instances

Figure 12: Execution time as a function of the amount of noise in the dataset




Figure 13: The memory usage as a function of the size of the dataset

1001 execution time of the J4.8 algorithm decreases as the amount of noise in the
 1002 dataset increases. It can be derived that the execution time decreases faster
 1003 when the percentage of noise is higher than 5%. As shown in the previous
 1004 section, the relative error quickly increases once the amount of noise rises
 1005 above 5%. This is because the J4.8 algorithm will more quickly decide that
 1006 it is no longer useful to try to split up the decision tree. On the one hand,
 1007 this leads to a lower detection rate as not all the criteria are discovered. On
 1008 the other hand, this decreases the needed execution time of the algorithm.

1009 Figure 13 illustrates the memory usage of the framework as a function of
 1010 the size of the dataset. The fluctuating pattern of the graphs can be explained
 1011 by the memory that is consumed by the *Garbage Collector* in Java. However,
 1012 trend lines can clearly be discerned. It can be noted that the memory usage

is linear as a function of the amount of instances. Moreover, the total amount of consumed memory stays quite low, i.e., at most 80 MB. For the realistic dataset of 1,680 instances the memory usage is negligible, namely about 10 MB.

It can be concluded that a dataset of realistic size for the SIRS use case can be processed by any modern PC or server and no cloud-based solutions are needed to run the framework.

6. Conclusions

In this paper a self-learning, probabilistic, ontology-based framework was presented, which allows context-aware applications to adapt their behavior at run-time. The proposed framework consists of the following steps. First, an ontology-based context model with accompanying rule-based context-aware algorithms is used to capture the behavior of the user and the context in which it is exhibited. Historical information is then gathered by algorithms that identify missing or inaccurate knowledge in the context-aware platform. This historical information is filtered, cleaned and structured so that it can be used as input for data mining techniques. The results of these data mining techniques are then prioritized and filtered by associating probabilities, which express how reliable or accurate they are. These results and the associated probabilities are then integrated into the context model and dynamic algorithms. These probabilities clarify to the stakeholders that this new knowledge has not been confirmed by rigorous evaluation. Finally, these probabilities are adapted, i.e., in- or decreased, according to context and behavioral information gathered about the usage of the learned information.

1037 The pipeline architecture of the framework was presented and its imple-
1038 mentation was detailed. Finally, a representative use case was presented to
1039 illustrate the applicability of the framework, namely mining the reasons for
1040 patients' nurse call light use to automatically launch calls. More specifically,
1041 detecting SIRS as a reason for nurse calls was used as a realistic scenario to
1042 evaluate the correctness and performance of the proposed framework. It is
1043 shown that correct results are achieved when the dataset contains at least
1044 1,000 instances and the amount of noise is lower than 5%. The execution
1045 time and memory usage are also negligible for a realistic dataset, i.e., below
1046 100 ms and 10 MB.

1047 Future work will mainly focus on the development of more intricate moni-
1048 toring, probabilistic relevance and filter algorithms. Moreover, a prototype of
1049 the proposed framework will be deployed and evaluated in a real-life setting.

1050 **Acknowledgements**

1051 We would like to acknowledge that part of this research was supported by
1052 the iMinds Project ACCIO co-funded by the IWT, iMinds and the following
1053 partners: Televic NV, Boone NV, Dominiek Savio Instituut and In-Ham.
1054 F. Ongenae and M. Claeys would like to thank the IWT for financial support
1055 through their Ph.D. grant. The authors thank F. De Backere for support in
1056 constructing the figures.

1057 **References**

1058 Anderston, J., Aydin, C., 1997. Evaluating the impact of health care infor-
1059 mation systems. International Journal Technology Assessment in Health

1060 Care 13 (2), 380–393.

1061 Baader, F., Calvanese, D., McGuinness, D. L., Nardi, D., Patel-Schneider,
1062 P., 2003. The Description Logic Handbook: Theory, Implementation and
1063 Applications. Cambridge University Press.

1064 Baldauf, M., Dustdar, S., Rosenberg, F., 2007. A survey on context-aware
1065 systems. *International Journal of Ad Hoc and Ubiquitous Computing* 2 (4),
1066 263–277.

1067 Baralis, E., Cagliero, L., Cerquitelli, T., Garza, P., Marchetti, M., 2011.
1068 Cas-mine: providing personalized services in context-aware applications
1069 by means of generalized rules. *Knowledge and Information Systems* 28 (2),
1070 283–310.

1071 Bardram, J., 14-17 March 2004. Applications of context-aware computing
1072 in hospital work - Examples and design principles. In: *Proceedings of the*
1073 *Annual ACM Symposium on Applied Computing*. New York, NY, USA:
1074 ACM Press;, Nicosia, Cyprus, pp. 1574–1579.

1075 Bass, L., Clements, P., Kazman, R., 2003. *Software Architecture in Practice*,
1076 2nd Edition. Addison-Wesley Professional.

1077 Blake, J. A., Harris, M. A., 2008. The Gene Ontology (GO) project: struc-
1078 tured vocabularies for molecular biology and their application to genome
1079 and expression analysis. *Current Protocols in Bioinformatics* 23 (7.2.1-
1080 7.2.9), 1472–6947, <http://www.geneontology.org/>.

1081 Bricon-Souf, N., Newman, C. R., 2007. Context awareness in health care: A
1082 review. *International Journal of Medical Informatics* 76 (1), 2–12.

1083 Burgelman, J.-C., Punie, Y., 2006. Close encounters of a different kind: am-
 1084 bient intelligence in Europe. *True Vision: The Emergence of Ambient In-*
 1085 *telligence*, 19–35.

1086 Byun, H., Cheverst, K., 2004. Utilizing context history to provide dynamic
 1087 adaptations. *Applied Artificial Intelligence* 18 (6), 533–548.

1088 Carroll, J. J., Dickinson, I., Dollin, C., Reynolds, D., Seaborne, A., Wilkin-
 1089 son, K., May 17–22 2004. Jena: implementing the semantic web recom-
 1090 mendations. In: *Proceedings of the 13th international conference on World*
 1091 *Wide Web, Alternate track papers & posters*. New York, NY, USA, pp.
 1092 74–83.

1093 Chen, H., 2004. An intelligent broker architecture for pervasive context-aware
 1094 systems. Ph.D. thesis, University of Maryland, Baltimore County.

1095 Chin, T., 2004. Technology valued, but implementing it into practice is
 1096 slow. *American Medical News*, [http://www.ama-assn.org/amednews/](http://www.ama-assn.org/amednews/2004/01/19/bisb0119.htm)
 1097 [2004/01/19/bisb0119.htm](http://www.ama-assn.org/amednews/2004/01/19/bisb0119.htm).

1098 Colpaert, K., Belleghem, S. V., Benoit, D., Steurbaut, K., Turck, F. D.,
 1099 Decruyenaere, J., October 11-14 2009. Has information technology finally
 1100 been adopted in intensive care units? In: *22nd Annual Congress of the*
 1101 *European Society of Intensive Care Medicine*. Vienna, Austria, p. 235.

1102 Criel, J., Claeys, L., 2008. A transdisciplinary study design on context-aware
 1103 applications and environments. a critical view on user participation within
 1104 calm computing. *Observatorio* 2 (2), 57–77.

- 1105 Davies, M. G., Hagen, P. O., 1997. Systematic inflammatory response syn-
1106 drome. *The British Journal of Surgery* 84 (7), 920–935.
- 1107 de Toledo, P., Jimenez, S., del Pozo, F., Roca, J., Alonso, A., Hernandez, C.,
1108 2006. Telemedicine experience for chronic care in copd. *IEEE Transactions*
1109 *on Information Technology in Biomedicine* 10 (3), 567–573.
- 1110 Dey, A. K., Abowd, G. D., 1-6 April 2000. Towards a better understanding
1111 of context and context-awareness. In: Morse, D. R., Dey, A. K. (Eds.),
1112 *Proceedings of the CHI Workshop on the What, Who, Where, When and*
1113 *How of Context-Awareness*. New York, NY, USA: ACM Press, The Hague,
1114 The Netherlands, pp. 304–307.
- 1115 Everitt, B. S., Skrondal, A., 2010. *The Cambridge Dictionary of Statistics*,
1116 4th Edition. Cambridge University Press, New Yor, USA.
- 1117 Fishkin, K., Wang, M., Fishkin, K. P., Wang, M., 2003. A flexible, low-
1118 overhead ubiquitous system for medication monitoring. Tech. rep., Intel
1119 Research Seattle, Technical Memo IRS-TR-03-011.
- 1120 Floerkemeier, C., Siegemund, F., 25 October 2003. Improving the effective-
1121 ness of medical treatment with pervasive computing technologies. In: *Pro-*
1122 *ceedings of the 2nd International Workshop on Ubiquitous Computing for*
1123 *Pervasive Healthcare Applications at International Conference on Ubiqui-*
1124 *tous Intelligence and Computing*. Seattle, Washington, USA.
- 1125 Fook, V. F. S., Tay, S. C., Jayachandran, M., Biswas, J., Zhang, D., 13-17
1126 March 2006. An ontology-based context model in monitoring and handling
1127 agitation behaviour for persons with dementia. In: *Proceedings of the 4th*

1128 IEEE International Conference on Pervasive Computing and Commununi-
1129 cations Workshops (PERCOMW). Washington, DC, USA: IEEE Com-
1130 puter Society;, Pisa, Italy, pp. 560–564.

1131 Gruber, T., 1993. A translation approach to portable ontology specifications.
1132 Knowledge Acquisition 5 (2), 199–220.

1133 Gu, T., Pung, H. K., Zhang, D. Q., 2005. A service-oriented middleware for
1134 building context-aware services. Journal of Network and Computer Appli-
1135 cations (JNCA) 28 (1), 1–18.

1136 Hong, J., Suh, E., Kim, S., 2009a. Context-aware systems: A literature
1137 review and classification. Expert Systems with Applications 36 (4), 8509–
1138 8522.

1139 Hong, J., Suh, E.-H., Kim, J., Kim, S., 2009b. Context-aware system for
1140 proactive personalized service based on context history. Expert Systems
1141 with Applications 36 (4), 7448–7457.

1142 Horrocks, I., Patel-Schneider, P. F., Boley, H., Tabet, S., Grosz, B., Dean,
1143 M., 2004. SWRL: A semantic web rule language combining OWL and
1144 RuleML. Tech. rep., <http://www.w3.org/Submission/SWRL/>.

1145 Hu, B., Hu, B., Wan, J., Dennis, M., Chen, H.-H., Li, L., Zhou, Q., 2010.
1146 Ontology-based ubiquitous monitoring and treatment against depression.
1147 Wireless Communications & Mobile Computing 10 (10), 1303–1319.

1148 Jahnke, J. H., Bychkov, Y., Dahlem, D., Kawasme, L., 2004. Context-aware
1149 information services for health care. In: Proc. of the Workshop on Modeling
1150 and Retrieval of Context. pp. 73–84.

- 1151 Jansen, B., Deklerck, R., November 29 2006. Context aware inactivity recog-
 1152 nition for visual fall detection. In: Proceedings of the Pervasive Health
 1153 Conference and Workshops. Innsbruck, Austria, pp. 1–4.
- 1154 Klinov, P., June 1-5 2008. Pronto: A non-monotonic probabilistic descrip-
 1155 tion logic reasoner. In: Proceedings of the 5th European Semantic Web
 1156 Conference. Tenerife, Spain, pp. 822–826.
- 1157 Korhonen, I., Paavilainen, P., Särelä, A., 25 October 2003. Application of
 1158 ubiquitous computing technologies for support of independent living of
 1159 the elderly in real life settings. In: Proceedings of the 2nd International
 1160 Workshop on Ubiquitous Computing for Pervasive Healthcare Applications
 1161 at International Conference on Ubiquitous Intelligence and Computing.
 1162 Seattle, Washington, USA.
- 1163 Kotsiantis, S. B., 2007. Supervised machine learning: A review of classifica-
 1164 tion techniques. *Informatica* 31 (3), 249–268.
- 1165 Kotsiantis, S. B., 2013. Decision trees: a recent overview. *Artificial Intelli-*
 1166 *gence Review* 39 (4), 261–283.
- 1167 Lukasiewicz, T., 2007. Probabilistic description logics for the semantic web.
 1168 Tech. rep., Technical University of Wien, Institute for Information Sys-
 1169 tems, Wien, Austria.
- 1170 McGuinness, D. L., Harmelen, F. v., February 10 2004. Owl web ontology
 1171 language overview. Tech. Rep. REC-owl-features-20040210, World Wide
 1172 Web Consortium, <http://www.w3.org/TR/owl-features/>.

- 1173 Mihailidis, A., Carmichael, B., Boger, J., Fernie, G., 25 October 2003. An
1174 intelligent environment to support aging-in-place, safety, and independence
1175 of older adults with dementia. In: Proceedings of the 2nd International
1176 Workshop on Ubiquitous Computing for Pervasive Healthcare Applications
1177 at International Conference on Ubiquitous Intelligence and Computing.
1178 Seattle, Washington, USA.
- 1179 Mitchell, S., Spiteri, M. D., Bates, J., Coulouris, G., 17-20 September 2000.
1180 Context-aware multimedia computing in the intelligent hospital. In: Pro-
1181 ceedings of the 9th workshop on ACM SIGOPS European workshop: be-
1182 yond the PC: new challenges for the operating system. New York, NY,
1183 USA: ACM;, Kolding, Denmark, pp. 13–18.
- 1184 Munoz, M. A., Rodríguez, M., Favela, J., Martinez-Garcia, A. I., González,
1185 V. M., 2003. Context-aware mobile communication in hospitals. *Computer*
1186 36 (9), 38–46.
- 1187 Nyström, P. O., 1998. The systemic inflammatory response syndrome: defi-
1188 nitions and aetiology. *Journal of Antimicrobial Chemotherapy* 41, 1–7.
- 1189 O'Connor, M. J., Das, A. K., 2010. A lightweight model for representing and
1190 reasoning with temporal information in biomedical ontologies. In: Interna-
1191 tional Conference on Health Informatics (HEALTHINF). Valencia, Spain,
1192 pp. 90–97.
- 1193 Ongenaes, F., Bleumers, L., Sulmon, N., Verstraete, M., Jacobs, A., Van Gils,
1194 M., Ackaert, A., De Zutter, S., Verhoeve, P., De Turck, F., 26-29 Octo-
1195 ber 2011a. Participatory design of a continuous care ontology: Towards

1196 a user-driven ontology engineering methodology. In: Filipe, J., Dietz, J.
1197 L. G. (Eds.), Proceedings of the International Conference on Knowledge
1198 Engineering and Ontology Development (KEOD). ScitePress Digital Li-
1199 brary;; Paris, France, pp. 81–90.

1200 Ongenaë, F., De Backere, F., Steurbaut, K., Colpaert, K., Kerckhove, W.,
1201 Decruyenaere, J., De Turck, F., 2011b. Appendix B: overview of the exist-
1202 ing medical and natural language ontologies which can be used to support
1203 the trans-lation process. BMC Medical Informatics and Decision Making
1204 10 (3), 4.

1205 Ongenaë, F., De Backere, F., Steurbaut, K., Colpaert, K., Kerckhove, W.,
1206 Decruyenaere, J., De Turck, F., 2011c. Appendix B: overview of the exist-
1207 ing medical and natural language ontologies which can be used to support
1208 the translation process. BMC Medical Informatics and Decision Making
1209 10 (3), 4.

1210 Ongenaë, F., Myny, D., Dhaene, T., Defloor, T., Van Goubergen, D., Verho-
1211 eve, P., Decruyenaere, J., De Turck, F., 2011d. An ontology-based nurse
1212 call management system (oNCS) with probabilistic priority assessment.
1213 BMC Health Services Research 11 (26), 28.

1214 Orwat, C., Graefe, A., Faulwasser, T., 2008. Towards pervasive computing
1215 in health care - a literature review. BMC Medical Informatics and Decsion
1216 Making 8 (26), 18.

1217 Paganelli, F., Giuli, D., 2011. An ontology-based system for context-aware

1218 and configurable services to support home-based continuous care. IEEE
1219 Transactions on Information Technology in Biomedicine 15 (2), 324–333.

1220 Prentzas, J., Hatzilygeroudis, I., 2007. Categorizing approaches combining
1221 rule-based and case-based reasoning. Expert Systems 24 (2), 97–122.

1222 Prud’hommeaux, E., Seaborne, A., January 15 2008. Sparql query language
1223 for rdf. W3C Recommendation REC-rdf-sparql-query-20080115, [http://](http://www.w3.org/TR/rdf-sparql-query/)
1224 www.w3.org/TR/rdf-sparql-query/.

1225 Quinlan, J. R., 1986. Induction of decision trees. Machine Learning 1 (1),
1226 81–106.

1227 Quinlan, J. R., 1993. C4.5: Programs for Machine Learning. Morgan Kauf-
1228 mann, San Francisco, CA, USA.

1229 Quinlan, J. R., 1996. Improved use of continuous attributes in c4.5. Journal
1230 of Artificial Intelligence Research 4 (1), 77–90.

1231 Rector, A. L., Rogers, J. E., Zanstra, P. E., van der Haring, E., 8-12 Novem-
1232 ber 2003. OpenGALEN: Open source medical terminology and tools.
1233 In: Proceedings of the annual American Medical Informatics Association
1234 (AMIA) Symposium. American Medical Informatics Association (AMIA);,
1235 Washington, DC, USA, p. 982, <http://www.opengalen.org/>.

1236 Rodríguez, M. D., Tentori, M., Favela, J., Saldaña, D., García, J.-P., 7-
1237 8 August 2011. Care: An ontology for representing context of activity-
1238 aware healthcare environments. In: Proceedings of the AAAI Workshop
1239 on Activity Context Representation: Techniques and Languages. Menlo
1240 Park, USA: AAAI Press;, Paris, San Francisco, CA, USA.

- 1241 Román, M., Hess, C., Cerqueira, R., Campbell, R. H., Nahrstedt, K., 2002.
1242 Gaia: A middleware infrastructure to enable active spaces. *IEEE Pervasive*
1243 *Computing* 1, 74–83.
- 1244 Rosse, C., Jr, J. L. V. M., 2008. The foundational model of anatomy ontology.
1245 In: Burger, A., Davidson, D., Baldock, R. (Eds.), *Anatomy Ontologies for*
1246 *Bioinformatics: Principles and Practice*. Springer;, London, UK, pp. 59–
1247 117.
- 1248 Russell, S., Norvig, P., 2003. *Artificial Intelligence, A Modern Approach*, 2nd
1249 Edition. Prentice Hall.
- 1250 Santos, L. O. B. S., Wijnen, R. P., Vink, P., 26-30 November 2007. A service-
1251 oriented middleware for context-aware applications. In: *Proceedings of*
1252 *the 5th International Workshop on Middleware for Pervasive and Ad hoc*
1253 *Computing*. New York, NY, USA: ACM Press;, Newport Beach, Orange
1254 County, CA, USA, pp. 37–42.
- 1255 Sirin, E., Parsia, B., Grau, B. C., Kalyanpur, A., Katz, Y., 2007. Pellet: A
1256 practical owl-dl reasoner. *Journal of Web Semantics* 5 (2), 51–53.
- 1257 Skov, M. B., Hoegh, R. T., 2006. Supporting information access in a hospital
1258 ward by a context-aware mobile electronic patient record. *Personal and*
1259 *Ubiquitous Computing* 10 (4), 205–214.
- 1260 Stanford, V., 2003. Beam me up, doctor mccoy. *IEEE Pervasive Computing*
1261 2 (3), 13–18.
- 1262 Strang, T., Linnhoff-Popien, C., 7 September 2004. A context modeling sur-
1263 vey. In: Indulska, J., Roure, D. D. (Eds.), *Proceedings of the 6th Inter-*

1264 national Conference on Ubiquitous Computing (UbiComp), Workshop on
1265 Advanced Context Modelling, Reasoning and Management. Nottingham,
1266 UK, pp. 31–41.

1267 Strobbe, M., Hollez, J., Jans, G. D., Laere, O. V., Nelis, J., Turck, F. D.,
1268 Dhoedt, B., Demeester, P., Janssens, N., Pollet, T., May 25 2007. De-
1269 sign of casp: an open enabling platform for context aware office and city
1270 services. In: Pfeifer, T., Strassner, J., Dobson, S. (Eds.), Proceedings of
1271 the 4th International Workshop on Managing Ubiquitous Communications
1272 and Services (MUCS 2007). Munich, Germany, pp. 123–142.

1273 Strobbe, M., Laere, O. V., Dhoedt, B., Turck, F. D., Demeester, P.,
1274 2012a. Hybrid reasoning technique for improving context-aware applica-
1275 tions. Knowledge and Information Systems 31 (3), 581–616.

1276 Strobbe, M., Laere, O. V., Ongenae, F., Dauwe, S., Dhoedt, B., Turck, F. D.,
1277 Demeester, P., Luyten, K., 2012b. Novel applications integrate location
1278 and context information. IEEE PERVASIVE COMPUTING 11 (2), 64–
1279 73.

1280 Su, J., Zhang, H., 2006. A fast decision tree learning algorithm. In: Pro-
1281 ceedings of the 21st National Conference on Artificial Intelligence. Boston,
1282 MA, USA, pp. 500–505.

1283 Suzuki, T., Doi, M., 31 March - 5 April 2001. Lifeminder: an evidence-based
1284 wearable healthcare assistant. In: Beaudouin-Lafon, M., Jacob, R. J. K.
1285 (Eds.), Proceedings of the ACM Conference on Human Factors in Com-

1286 puting Systems. New York, NY, USA: ACM, Seattle, Washington, USA,
1287 pp. 127–128.

1288 Tentori, M., Segura, D., Favela, J., 2009. Chapter VIII: Monitoring hospi-
1289 tal patients using ambient displays. In: Olla, P., Tan, J. (Eds.), Mobile
1290 Health Solutions for Biomedical Applications, 1st Edition. Vol. 1. Medical
1291 Information Science Reference;, Hershey, New York, USA, pp. 143–158.

1292 Tsang, S. L., Clarke, S., 2008. Mining user models for effective adaptation
1293 of context-aware applications. *International Journal of Security and its*
1294 *Applications* 2 (1), 53–62.

1295 Valls, A., Gibert, K., Sánchez, D., , Bateta, M., 2010. Using ontologies for
1296 structuring organizational knowledge in home care assistance. *Internation*
1297 *Journal of Medical Informatics* 79 (5), 370–387.

1298 Varshney, U., 2009. Chapter 11: Context-awareness in healthcare. In: *Perva-*
1299 *sive Healthcare Computing: EMR/EHR, Wireless and Health Monitoring*,
1300 1st Edition. Springer Science + Business Media, LLC;, New York, NY,
1301 USA, pp. 231–257.

1302 Verstichel, S., De Poorter, E., De Pauw, T., Becue, P., Volckaert, B., De
1303 Turck, F., Moerman, I., Demeester, P., 2010. Distributed ontology-based
1304 monitoring on the IBBT WiLab.t infrastructure. In: *Proceedings of the*
1305 6th International Conference on Testbeds and Research Infrastructures
1306 for the Development of Networks and Communities (TridentCom). Berlin,
1307 Germany, pp. 509–525.

- 1308 Witten, I. H., Frank, E., Hall, M., 2011. Data Mining: Practical Machine
1309 Learning Tools and Techniques, 3rd Edition. Morgan-Kaufmann.
- 1310 Xue, W., Pung, H. K., 2012. Chapter 8: Context-Aware middleware for
1311 supporting mobile applications and services. In: Kamur, A., Xie, B.
1312 (Eds.), Handbook of Mobile Systems Applications and Services, 1st Edi-
1313 tion. Vol. 1. CRC Press;, Florida, USA, pp. 269–304.
- 1314 Yilmaz, O., Erdur, R. C., 2012. iconawa - an intelligent context-aware system.
1315 Expert Systems with Applications 39 (3), 2907–2918.
- 1316 Zhang, D., Yu, Z., Chin, C.-Y., 2005. Context-aware infrastructure for per-
1317 sonalized healthcare. Studies in Health Technology and Informatics 117,
1318 154–163.