Knowledge Engineering Requirements for Generic Diagnostic Systems

Andreas Mueller[1], Ingmar Hofmann[1], Heiner Oberkampf[2] and Sonja Zillner[2]
[1]*Siemens AG, Industry Sector, Advanced Technologies & Standards, Nuremberg, Germany*
[2]*Siemens AG, Corporate Technology, Munich, Germany*

Keywords: Diagnostic Knowledge, Knowledge Representation, Knowledge-based Diagnostics, Domain-specific Language, Requirements Analysis, Knowledge Engineering, Reasoning, Knowledge Interchange.

Abstract: Diagnostics is the process of determining the nature of malfunctions or faults of systems in various domains. With regard to the complexity of systems and their composition of different subsystems or subcomponents, for which different diagnostic approaches are optimal, no means exist for seamless and agile cooperation and information exchange between currently isolated diagnostic approaches. However, we consider this essential for an integrated diagnostic mechanism covering complex systems in their entirety. Hence, in this paper, we show the basic requirements for a generic diagnostic knowledge representation language (DKRL) by investigating typical diagnostic examples from different domains, namely industry and medicine. DKRL is intended to facilitate the generic representation, handling, and interchange of diagnostic knowledge required for performing diagnostics without regard to specific diagnostic approaches.

1 INTRODUCTION

We are concerned with the efficient representation of diagnostic information.

Being the process of analytically determining the nature and cause(s) of a malfunction of a technical or biological system using structural, functional and causal knowledge about the system, diagnostics is used in various disciplines, with variation in the use and representation of the underlying knowledge. On malfunctions, a useful diagnostic mechanism must be able to determine defective parts of the system and also the root cause of these failures, which may sometimes lie outside the system's boundaries. This determination relies on observations delivering more details on the system's current behavior.

Various diagnostic mechanisms address isolated diagnostic problems. However, the seamless and agile cooperation between them is currently not possible. A truly integrated diagnostic mechanism for analyzing complex systems requires a high-level exchange of diagnostic information across isolated diagnostic mechanisms. For this, a vast amount of different information has to be managed in a precise and standardized manner. These challenges can be met by a generic diagnostic representation language providing means for handling diagnostic knowledge.

Current computer-based diagnostic systems use established reasoning methods, e.g. rule-based reasoning (Ligeza, 2006), case-based reasoning (Kolodner, 1993) or probabilistic reasoning (Pearl, 2005). However, there is little common ground regarding the formalisms for representing the diagnostic knowledge. In fact, these are tightly coupled to each approach. A good survey regarding such knowledge representation is provided in (Van Harmelen, Lifschitz and Porter, 2008). This, however, causes a semantic gap between the perception of a diagnostic problem and its formal representation. Also, this causes the diagnostic knowledge to be bound to the employed diagnostic approach. The flexible exchange of information about the structure of the system being diagnosed, about malfunctions, and observations requires an explicitly specified terminology of a generic diagnostic process in an unambiguous manner.

Our overall research goal is to develop a generic diagnostic knowledge representation language (which will be referred to as DKRL) that addresses this lack of a coherent diagnostic formalism in the form of a domain-specific language (DSL). Intended to facilitate the representation of diagnostic problems as such with all relevant diagnostic aspects, but independent of any reasoning approach.

Mueller A., Hofmann I., Oberkampf H. and Zillner S..
Knowledge Engineering Requirements for Generic Diagnostic Systems.
DOI: 10.5220/0004124601840189
In *Proceedings of the International Conference on Knowledge Engineering and Ontology Development* (KEOD-2012), pages 184-189
ISBN: 978-989-8565-30-3

DKRL could eventually become the basis for machine-processable diagnostic lexicons for arbitrary domains. Hence, we regard two aspects as being equally important: representational capabilities for the diagnostic knowledge itself as well as facilities for the efficient handling of represented knowledge. This paper discusses the analysis and gathering of requirements that need to be addressed by such a DSL for diagnostics.

The paper is organized as follows. Section 2 describes previous work. Section 3 illustrates aspects of diagnostics in different domains. Section 4 investigates diagnostics in the domains and introduces the requirements for generic knowledge representation. Section 5 concludes the paper and provides an outlook.

2 RELATED WORK

As one basic distinction of diagnostic systems, we have *model-based* (Lucas, 1998) or *first principles* (Reiter, 1987) diagnosis, and *heuristic classification* (Lucas, 1998) or *heuristic diagnosis* (Reiter, 1987). Model-based diagnosis as consistency-based diagnosis and abductive diagnosis (Lucas, 1998)(Poole, 1994) mainly proved useful in the technical/industrial domain, whereas in the medical domain, heuristic diagnosis is often used. In model-based diagnosis we have a description about how the system is meant to operate, together with observations. In heuristic diagnosis, information like “rules of thumb, statistical intuition and past experience” are more important and “the real world system being diagnosed is only weakly represented” (Reiter, 1987). Even in model-based diagnosis, there are many different formalisms for similar problems.

Existing approaches for a generic representation language for diagnostic knowledge focus on only one of the diagnosis problems. (Reiter, 1987) and (Poole, 1994) focus on model-based diagnosis. In (Poole, 1994), a further distinction of system-driven diagnosis in Consistency-Based Diagnosis and Abductive Diagnosis is made. It is shown that for a certain class of problems both formalisms reach the same diagnosis. In (Lucas, 1998) an attempt is made to create a generic diagnosis language. “Evidence functions” are used to represent the knowledge common to all diagnostic systems, the interactions among defects and findings (Lucas, 1998). The experience-driven (heuristic) approach is realized in Bayesian networks (probabilistic dependencies), default logic (rules of thumb) etc. However, there is still no overall diagnosis representation language able to represent the full spectrum of different diagnostic knowledge.

As our overall goal, DKRL is intended to cover both model-based and heuristics-based diagnosis. Showing typical features of the respective diagnosis types, in the following the industrial and the medical domain were selected for a requirements analysis.

3 DIAGNOSTICS IN DIFFERENT DOMAINS

We exemplarily consider the domains “industry” and “medicine” since these substantially differ in complexity and availability of reliable factual and causal knowledge, yet in both domains reaching a correct diagnosis quickly is critical.

The proposal to capture the notion of diagnostic reasoning has been considered by two extreme poles of the diagnosis problem (Poole, 1994): Firstly, the overall aim may be to describe how components are structured and work normally, however information on the origin and the manifestation of malfunctions is missing. This holds true for the industrial domain, thus, diagnostic algorithms aim to isolate deviations from normal behavior. Secondly, knowledge about faults and symptoms may be used to interpret the relevance of abnormalities. This holds true for the medical domain: medical diagnostic knowledge is typically about “incorrect functioning”.

For a comprehensive set of requirements needed to represent diagnostic knowledge generically, we analyze typical diagnostic use case scenarios from the industrial and the medical domains. The following examples illustrate the aspects relevant in diagnostic processes in the selected domains and show the requirements to be met in order to perform the described diagnostics. When gathering the requirements, we discussed with experts and analyzed existing systems to identify roadblocks and shortcomings. The medical examples are taken from interviews with our clinical partners.

3.1 Diagnostics in Industry

Typically, the industry domain shows a high degree of engineered knowledge, with an adequate understanding of the considered plant or component and corresponding diagnostic knowledge being possibly available from the beginning of the respective lifecycle. Thus, observations can often be performed directly and symptoms can often be treated as directly identifiable causes of observed

Table 1: Entities from the industrial domain.

#	Entity	Explaining remark
I1.	Causal relationships	In many cases it is possible to state the potential causes for symptoms.
I2.	Causes and effects (symptoms)	Sensors as well as the human operator's sense provide information that can be interpreted as a manifestation of a malfunction or fault.
I3.	Context-dependent interpretation	Nominal values for operating parameters need to be interpreted with regard to e.g. the currently selected mode of plant operation or environmental influences.
I4.	Likelihood of occurrence	As shown in the example of the feeder malfunction, of two possible causes one might occur less likely than the other.
I5.	Localization	Recognizing components or functions of a plant where a given effect typically occurs is important in order to direct service technicians.
I6.	Probabilistic causality	Causal relations are not necessarily absolute. Instead, there would be more than one possible, but not equally likely, causes for an observed symptom..
I7.	Significance of a symptom	Certain observed symptoms are typical for certain causes, due to the respective system's structure or functionality.
I8.	Temporal correlation	Some causes and effects become relevant only in correlation to the amount of time passed. Also, due to aspects such as the system's structure or functionality, the effect of an occurred cause might become visible only after some time has passed.
Additionally important functional aspects		
I9.	Extensibility	Knowledge is subject to change due to plant modifications during the lifecycle.
I10.	Incomplete knowledge	A lack of knowledge at the system level and instance level must be handled. Exact probabilistic knowledge about the causal relations is not always available.
I11.	Reusability	Knowledge about symptoms, etc. may be relevant for different faults, thus should be reused.

faults by applying sensors to critical positions in the structure or the process. Also, the causality behind symptoms is usually rather easy to determine.

3.1.1 Example 1: Process Industry

(Abdul-Wahab et al., 2007) give examples of troubleshooting in the domain of multi-stage flash seawater desalination, with focus on the brine heater component. Considering the process value of "condensate conductivity" as a representation of the salt ratio in low pressure steam after condensation, the plant operator might observe a gradual increase. The degree of increase might be a symptom of leaking tubes in the brine heater, since this would cause seawater to mix with the condensate, resulting in an increase of conductivity due to a higher salt ratio. Hence, the "leakage" needs to be ruled out by maintenance actions. If the conductivity continues to rise, automatic valve operation measures would be taken. Similarly, the conductivity of the distillate might also show a symptomatic, differently located increase. Here, a possible cause might be an increased "top brine temperature", which can be responsible for increased brine flashing, causing the conductivity to rise due to a higher salt ratio.

3.1.2 Example 2: Manufacturing Industry

At Siemens in Nuremberg a research facility is operated, producing a running text made from small plastic disks using a series of conveyor belts (see Figure 1) as an exemplary manufacturing process. Normally, the disks are separated from the rear-side storage belt and transported onto the right-hand-side "column belt" (1), where the next required disk column is prepared according to the control system. The column is then pushed onto the front-side "text belt" by a properly triggered proximity switch (PS)-controlled feeder (2), positioning the column at the correct distance in relation to the previous one. Considering the "text belt", the symptom of "irregular positioning patterns" might occur.

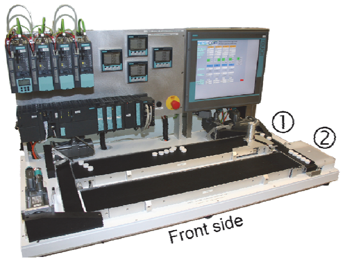

Figure 1: Siemens research facility for the manufacturing process of mechanical running text production.

This can be caused by a wrong feeder mode (continuous instead of on-demand). Currently, the reason may be either a broken PS cable connection, so trigger signals can no longer be received, or a broken PS, which can be treated as less likely.

3.2 Diagnostics in Medicine

In the medical domain, direct observations are often

Table 2: Entities from the medical domain.

#	Entity	Explaining remark
M1.	Context-dependent interpretation	Standard intervals for measurements like blood pressure or blood counts need patient-specific interpretation regarding, e.g. the patient's age.
M2.	Functional dependencies and localization	For instance, an artery functionally supplies blood to a number of organs. Such information allows the physician to anticipate phenomena and to imagine an internal situational picture of the disease he is confronted with.
M3.	Patient history	For many diseases risk factors are known. They influence the probability of certain diseases and allow the clinician to check for such diseases first.
M4.	Probabilistic causality	The disease-symptom-relation is typically a probabilistic causal relation – e.g. *lymphoma* has the associated symptom *enlarged lymph nodes* in 80% and the symptom *enlarged spleen* in 30% of all cases (Herold, 2011).
M5.	Relations between symptoms	Symptom descriptions are fuzzy, e.g. *elevated temperature* vs. *mild fever*. Also, some symptoms form groups like e.g. the B-symptoms - occurring together they constitute a cardinal symptom for *lymphoma*.
M6.	Symptom significance	Cardinal symptoms help to focus on certain diseases in the initial diagnosis phase. Additional occurrence of pathognomonic symptoms allow an immediate diagnosis.
M7.	Symptoms	There are symptoms that can be observed on a patient during a diagnostic process without being able to immediately draw conclusions about the actual disease. This situation is typical when recording a patient's medical history, where symptoms are collected as reported by a patient before being interpreted.
M8.	Temporal progression	Diseases may show a characteristic development of symptoms over time. Another temporal factor is the "novelty" of a symptom.
M9.	Urgency of symptoms	Diseases and their symptoms are not equally dangerous for the patient's health. Thus, highly dangerous diseases and related symptoms have to be checked first.
Additionally important functional aspects		
M10.	Extensibility	Knowledge is subject to change, so it must be represented in an extendable way.
M11.	Incomplete knowledge	A lack of knowledge at the system level and instance level must be handled. Exact probabilistic knowledge about the causal relations between diseases and symptoms is not always available. Similarly, the complete patient situation is seldom known.
M12.	Reusability	Diagnostic knowledge may be relevant for various diseases, thus should be reused.

impossible, so conclusions about the actually desired biological characteristics have to be drawn based on observable characteristics and presumed or confirmed interrelations. Because of the domain's inherent complexity, medical diagnostic knowledge contains a high degree of uncertainty.

3.2.1 Example 1: Differential Diagnosis

Differential diagnosis is a standard diagnostic approach in clinical practice. We illustrate the process along an example: *first, the observation of an initial set of (unspecific) symptoms,* say "fever", "night sweats", "feeling weak" and "changes in bowel patterns", *leads the clinician to suspect a set of likely diseases which might have caused the symptoms. Second, the set of likely diseases is turned into a ranked list based on information about cardinal symptoms, incidence proportion of a disease and other factors.* Since "changes in bowel patterns" is a cardinal symptom for "diverticulitis" and "colorectal cancer", we may obtain a ranked list like "diverticulitis", "colorectal cancer", "cold", "lymphoma" etc. *Third, the clinician aims to differentiate between the likely diseases by checking for further symptoms that might strengthen or weaken diagnoses on the list. At first, he will check other (cardinal) symptoms of top-ranked diseases.* In this case he might identify "weight loss" and "enlarged lymph nodes" but does not find evidence for "blood in stool" and "thickened intestinal wall". *Fourth a more precise list of likely diseases is obtained,* with "lymphoma" now at the top as both cardinal symptoms ("enlarged lymph nodes") and B symptoms (correlation of "fever", "night sweats", and "weight loss") are present. "Diverticulitis" and "colorectal cancer" become less likely as important symptoms "blood in stool" and "thickened intestinal wall" are absent. *The process continues until a plausible diagnosis is found.*

3.2.2 Example 2: Lyme Borreliosis

As an infectious disease, Lyme borreliosis has exactly one cause: the patient has been infected by bacteria of genus Borrellia after a tick bite (Masuhr, 1996). Depending on both the part of the body the infection took place at, and the time passed since the infection, different symptoms occur. The course of borrelliosis is divided into three stages after

Table 3: Requirements and classifications.

#	Requirement	Explaining remark	Based on
Core requirements: mandatory for the intended representation language			
R1.	Causality	DKRL must represent causality (cause-effect-relationships).	I1, M4
R2.	Causes and effects	DKRL must represent causes and the effects of causes.	I2, M7
R3.	Context-dependent interpretations	DKRL must represent information about the interpretation of measurements taking into consideration DO-specific influences.	I3, M1
R4.	Faults	DKRL must represent faults that may (have) occur(ed) on a DO.	I2, M3
R5.	Likeliness under error conditions	DKRL must represent the likeliness of symptoms and faults for each of the DO's components as well as for the DO as a whole to occur under error conditions.	I4, M6
R6.	Likeliness under nominal operating conditions	DKRL must represent the probability that symptoms and faults occur under nominal operating conditions, since even under optimum conditions there is possibility of spontaneous malfunctions that might spread throughout the DO.	I6, M4
R7.	Localization (physical)	DKRL must represent where a symptom or fault is located physically .	I5, M2
R8.	Localization (functional)	DKRL must represent where a symptom or fault is located functionally.	I5, M2
R9.	Significance of causal relationships	DKRL must represent that a symptom or cause may have a different significance to a fault or an effect, respectively, than another symptom/cause.	I6, I7, M4
R10.	Symptoms	DKRL must represent symptoms that may occur on a DO.	I2, M7
R11.	Temporal classification	DKRL must represent information that effects might become observable only after some time has passed after the occurrence of a cause.	I8, M8
Application-specific requirements: optional for the intended representation language			
R12.	Novelty of symptoms	DKRL must represent information that the new occurrence of symptoms during the course of a fault may be of special importance.	M8
R13.	Relationships between symptoms/causes/effects	DKRL must represent information that certain symptoms, causes, or effects have a special relationship to other symptoms, causes, or effects, respectively.	M5
R14.	Temporal relevance of symptoms	DKRL must represent the temporal relevance of symptoms, i.e. the period of time a symptom is of significance for a certain fault.	M8
R15.	Urgency	DKRL must represent information that certain symptoms need to be investigated prior to others.	M9
Overall functional requirements: mandatory for effective and efficient knowledge management			
R16.	Extensibility	DKRL must represent knowledge in an easily modifiable manner.	I9, M10
R17.	Incompleteness of knowledge	DKRL must represent knowledge so that valid descriptions can be created, even though there might be information missing.	I10, M11
R18.	Reusability	DKRL must represent knowledge in a manner that allows reuse or referencing of knowledge once it has been captured.	I11, M12

infection. The cardinal symptom of the first stage (3 days to 3 weeks) is a circular rash called erythema chronicum migrans at the region of the tick bite. In the second stage (1 to 4 months) different body parts will be affected and the patient might show symptoms of a meningopolyneuritis like radicular pain or even facial paralysis. Here, for anamnesis the physician would ask the patient if he remembers a tick bite, but also needs to conduct several tests, since these symptoms might be caused by other diseases. In the third stage (>5 months) untreated patients might show colored areas of skin (acrodermatitis chronica atrophicans) or joint disorders as a possible symptom of Lyme arthritis and diagnosis is even more complex. The common underlying reason for this diversity of symptoms is the spreading of the bacteria in the body together with the initiated pathomechanism (comprising the effects of "cellular invasion" and consequently "inflammatory reactions"), which progresses at different rates depending on the affected tissues.

4 REQUIREMENTS ON KNOWLEDGE REPRESENTATION

From these examples, we obtain the aspects to be represented and derive the requirements for DKRL. We distinguish between *core requirements* (most basic and essential; inherent to all diagnostic decisions), *application-specific requirements* (only relevant within applications for special diagnostic problems), and *overall functional requirements* (for effective and efficient handling of knowledge) (see Figure 2).

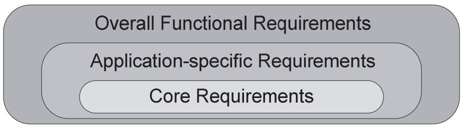

Figure 2: Requirement layers.

4.1 Relevant Entities from the Considered Domains

In the following sections we describe the knowledge entities and aspects typical for each domain. The representation language must be able to represent these either implicitly or explicitly.

4.1.1 Entities from the Industrial Domain

The examples illustrate that diagnostics in the industrial domain strongly uses pre-engineered domain knowledge. If available, systemic information about the representatives of the respective domains (i.e. components used in process technology or in the manufacture of discrete parts) in terms of structure and functionality is rather certain and complete. Hence, symptoms and often their causes are well-known and thus can be directly captured. To represent the required diagnostic knowledge, we have identified the entities in Table 1 to be required.

4.1.2 Entities from the Medical Domain

The examples illustrate that medical diagnosis is largely based on the clinician's experience and statistical information: he knows from experience which disease might cause certain general and cardinal symptoms, and he knows about the significance of symptoms for certain diseases. Based on statistics, more frequent diseases will be considered first. On the other hand, systemic information about the human organism is less certain and complete, and processes in the human body are highly interconnected and not yet understood well enough to facilitate correct model-based diagnosis. We consider the entities in Table 2 to be important in order to represent medical diagnostic knowledge.

4.2 Requirements and Classifications

From the entities listed for each domain, in Table 3 we now derive the requirements that the intended representation language has to meet (representatives from the domains are generically referred to as *diagnostic objects,* abbreviated *DO*).

5 CONCLUSIONS AND OUTLOOK

In this paper we have derived the requirements for a generic diagnostic knowledge representation language (DKRL) by investigating diagnostic knowledge from the exemplary domains of industry and medicine. DKRL is intended for use with any diagnostic system, for handling diagnostic knowledge in an easily reusable way. We have shown that the majority of requirements holds true in both domains. Application-specific requirements are mainly induced by the medical domain. Still, this does not restrict their relevance to the medical context. In fact, we consider that fulfilling these requirements step-by-step forms a suitable basis for gradual extension of the diagnostic functionalities addressable by DKRL. Hence, the identified requirements allow for the development of DKRL as well as a corresponding software infrastructure.

The future development of DKRL will focus on a prototypical implementation with full coverage of the core requirements and the overall functional requirements, followed by adding application-specific representational capabilities.

ACKNOWLEDGEMENTS

The authors are grateful to Dr. Kristina Bayerlein, senior physician at the University Hospital of Erlangen, for details on the interpretation of the medical terminology.

REFERENCES

Abdul-Wahab, S., Elkamelb, A., Al-Weshahic, M. & Al Yahmadia, A., 2007. Troubleshooting the brine heater of the MSF plant using a fuzzy logic-based expert system. In *Desalination*, vol. 217, no. 1-3, pp. 100-117.

Herold, G., 2011. *Innere Medizin*, Herold.

Kolodner, J., 1993. *Case-Based Reasoning,* Morgan Kaufmann Publishers Inc., San Mateo, CA.

Ligeza, A., 2006. *Logical Foundations for Rule-Based Systems,* Springer, Berlin Heidelberg.

Lucas, P., 1998. Analysis of notions of diagnosis. In *Artificial Intelligence*, vol. 105, no. 1-2, pp. 295-343.

Masuhr, K., Neumann, M., 1996. *Neurologie*, Hippokrates. Stuttgart.

Pearl, J., 2005. *Causality*, Cambridge University Press, New York.

Poole, D., 1994. Representing Diagnosis Knowledge. In *Annals of Mathematics and Artificial Intelligence*, vol. 11, pp. 33-50.

Reiter, R., 1987. A Theory of Diagnosis from First Principles. In *Artificial Intelligence*, vol. 32, no. 1, pp. 57-95.

Van Harmelen, F., Lifschitz, V., Porter, B., 2008. *Handbook of Knowledge Representation*, Elsevier Science.