A Hybrid Architecture for a Preoperative Decision Support System Using a Rule Engine and a Reasoner on a Clinical Ontology

Matt-Mouley Bouamrane[1,2], Alan Rector[1], and Martin Hurrell[2]

[1] School of Computer Science
Manchester University, UK
{mBouamrane,Rector}@cs.man.ac.uk
[2] CIS Informatics, Glasgow, UK
martin.hurrell@informatics.co.uk

Abstract. We report on a preventive care software system for preoperative risk assessment of patient undergoing elective surgery. The system combines a rule engine and a reasoner which uses a decision support ontology developed with a logic based knowledge representation formalism. We specifically discuss our experience of using a representation of a patient's medical history in OWL, combined with a reasoning tool to suggest appropriate preoperative tests based on an implementation of preoperative assessment guidelines. We illustrate the reasoning functionalities of the system with a number of practical examples.

1 Introduction

The primary goal of clinical guidelines is to improve the quality and efficiency of healthcare delivery, while providing an efficient, cost-effective and consistent service across healthcare institutions. Clinical guidelines are expressed as statements, rules, recommendations, management protocols, etc. They suggest a certain course of action given a specific medical *context*. Clinical guidance provide support to health professionals without overruling their clinical judgment and ability to use their own discretion in order to make decisions appropriate to the individual circumstances of the patients and the broader medical context. Clinical guidelines are generally designed after a lengthy and thorough process, involving clinical topic selection, consultation with stakeholders, establishing an expert panel which will examine the best evidence available on the topic, the generation of draft recommendations, typically by consensus, submission to stakeholders or peer-review, and finally, issuance of final guidelines. The guidelines may be used for training health professionals and supporting them in the intellectually demanding and knowledge-intensive environment of the health services. In addition, guidelines may lead to a significant improvement in service efficiency and a substantial reduction in costs. The study by Ferrando et al. on 702 patients undergoing preoperative assessment estimated that applying preoperative guidelines would reduce the cost of preoperative tests by 63% and in excess of

A. Polleres and T. Swift (Eds.): RR 2009, LNCS 5837, pp. 242–253, 2009.

40% for hospital stays per patient [1]. In addition, tests performed on patients, which were not deemed relevant by the guidelines, did not provided additional clinical information of interest.

While the benefits of clinical guidelines may be confirmed by a number of studies, the problem remains as ever of how to apply these effectively on the ground and integrate them in clinical workflows and work practices, amidst the huge volume of medical information available to health professionals, a proliferation of guidelines, as well as rapidly changing health policies [2]. The difficulty is thus to deliver the guidelines in a format which make them readily usable by health professionals. In addition, the considerable effort in producing and implementing the guidelines is another incentive to make these guidelines sharable and interoperable across services and institutions. Thus, the use of Information Technology in the health services, and in particular Clinical Decision Support Systems (CDSS), can potentially efficiently support health professionals in their duties. Tasks which may routinely be performed by CDSS include: efficiently managing patient clinical information, handling clinical guidelines, providing risk assessment, preventing errors of omission through alerts and reminders, extracting information of interest and making recommendations appropriate to the circumstances of the patients, such as treatment protocols [3,4,5]. Indeed, the potential to exploit computer systems to implement and share clinical guidelines has long been recognised, as evidenced by the development of the GuideLine Interchange Format (GLIF) [6], and efforts to develop guidelines models and guidelines based decision support systems [7,8].

We here report on our work combining a preventive care software system for preoperative risk assessment of patient using a rule engine with a decision support ontology developed with a logic based knowledge representation formalism. We specifically discuss our experience of using a representation of a patient's medical history in OWL, combined with a reasoning tool to suggest appropriate preoperative tests based on an implementation of the NICE preoperative assessment guidelines[1]. The paper is structured as follows: first we present the research motivation and background information on the preoperative decision support system. We provide an overview of the system decision support ontology and reasoning functionalities. We then discuss in details our implementation of the NICE preoperative assessment guidelines. We conclude with a discussion on this approach and directions for future work.

2 Background

This work is part of an ongoing project to introduce *semantic* technology into a preoperative risk assessment system software called Synopsis. The overall architecture of the preoperative assessment system is illustrated in Figure 1. The use of knowledge representation and reasoning both completes functionalities and overcomes a number of limitations of the existing system. While numeric risk

[1] NHS, National Institute for Clinical Excellence
http://www.nice.org.uk/Guidance/CG3

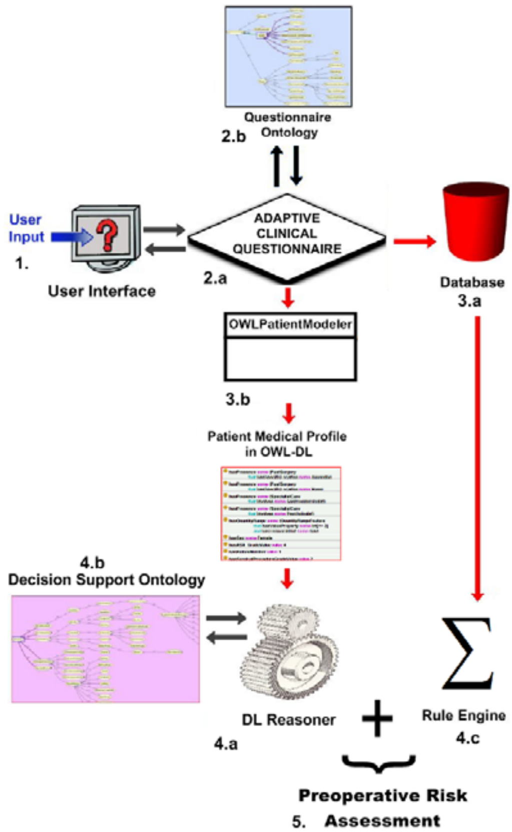

Fig. 1. Hybrid architecture of the rule-engine / clinical knowledge-base preoperative risk assessment system

score calculation is currently easily done by the system using an open source java-based rule engine, JBoss Rules[2], other tasks including categorisation, classification and logical inference were beyond the capacity of the system prior to introduction of semantic technology.

Functionalities introduced in the system as the result of this project include the development of an adaptive questionnaire, whereby clinical content of the information collection process adapts automatically to information specific to a patient clinical profile. Another enhancement to the system includes the ability to generate a high level semantic representation of a patient clinical history in the web ontology language OWL-DL. This is subsequently used to provide decision support functionalities via the use of automated reasoning tools. The web ontology language OWL is supported by an active research community, readily

[2] http://www.jboss.com/products/rules

and freely available development and reasoning tools and a well maintained application programming interface [9,10,11].

High level medical history representation is done automatically for new patients entered in the system and we have proposed a methodology for the semi-automatic generation of this medical history from legacy clinical databases in [12]. Note that prior to introduction to semantic technology to the system, the preoperative software was composed only of the following elements: user input (step 1), clinical data storage (3.a) and rule engine (4.c). Therefore, the preoperative risk assessment (5) was almost entirely based on the calculation of numeric scores. Thus, the introduction of semantic based technology in the system enables adaptive information collection (2a and 2b), high level semantic patient modelling (3.b) and decision support based on classification (4.a and 4.b) rather than numeric rules only. This provides for a significant enhancement to the functionalities and capabilities of the system. We refer the interested reader to [13,14,15,16] for a detailed description of the system features and functionalities.

3 Decision Support Ontology and Reasoning Functionalities

Once the system has completed the adaptive information collection phase and has generated a patient medical history semantic profile (3.b), reasoning on the decision support ontology can then be performed. The purpose of the decision support is to provide advice to clinicians and flag potential risks of complications so appropriate preventive measures can be taken accordingly. The decision support ontology is divided into three domains with distinct purposes. The first one consists of a Risk Assessment ontology and its purpose is to highlight potential intra-operative and post-operative complications given a patient medical profile and the planned surgical procedure. The second one is the Recommended Test ontology, which purpose is to suggest certain preoperative tests, which may help to decide whether it is safe to proceed with a planned surgery or whether further actions need to be taken (e.g. referral to a specialist, optimisation of patient's health prior to surgery, etc.) The Recommended Test ontology we are currently using is based on the preoperative tests clinical guidelines issued by the British National Health System (NHS) National Institute for Clinical Excellence (NICE). This will be discussed further in the next section. Finally, the last domain of the Decision Support ontology is the Precaution Protocol ontology, which may suggest a management or treatment protocol given a specific medical context.

In the system, decision support is usually provided in a 2 step process. The first step typically calculates risk scores or derives risk grades (ASA grades, surgical risk grades, etc.) using numerical formulas such as the Goldman and Detsky cardiac risk index, the Physiological and Operative Severity Score for the enUmeration of Mortality and Morbidity (POSSUM) [17], etc. Once the risk grades and categories have been derived from the first risk calculation step, the system can then perform decision support using the open-source java-based

PELLET reasoner [18] to reason on the decision support ontology given a patient OWL medical history profile.

4 Recommended Preoperative Investigations

4.1 Description of Investigation Guidelines

The purpose of the NICE guidelines recommendations for preoperative investigations is both to avoid patients undergoing unnecessary investigations, which can be detrimental to their health, as well as more efficiently managing limited resources in the public health services. The implementation of the NICE guidelines in the decision support system is both: (i) a pragmatic and useful functionality provided to health professionals and (ii) a good example of how the use of a clinical ontology and reasoner can provide functionalities beyond the capabilities of a traditional rule engine. Regarding the first point, the table in Figure 2 illustrates the format of the NICE guidelines recommendation for preoperative investigations. We here describe the guidelines in more details:

Type of investigations: the guidelines include 9 potential investigations: Chest X-Ray, ECG (Electrocardiogram), Full Blood Count, Haemostasis, Renal Function, Random glucose, Urine analysis, Blood gases and Lung Function.

Type of recommendations: there are currently 3 types of recommendations for each test: "test recommended", "test not recommended" and "consider test".

Factors Influencing recommendations: There are 5 factors taken into consideration in order to find the relevant recommendations: the (i) age of the patient, (ii) his ASA grade, (iii) the type of comorbidities the patient has (e.g. respiratory, cardiovascular, renal) (iv) the type of surgery (e.g. cardiovascular surgery, neurosurgery, etc.) (v) the risk grade of the surgery (from 1 to 4).

Number of cases in the guidelines: the guidelines are summarised for preoperative health assessors into 36 tables such as the one illustrated in Figure 2. There are different tables for different combinations of the 5 factors previously described, including different tables for children under 16 years old and adults over 16 years old. In total, there are at least 1242 different possible cases.

Perhaps not surprisingly, we found that in practice, preoperative health assessors faced considerable difficulties in using the guidelines. The important number of factors to take into consideration in order to find the correct table and then the specific case within this table, combined with the significant number of tables meant that too much time was being spent by preoperative health assessors trying to refer to the correct case. In addition, the preoperative health assessors would need to be able to categorise (i) the type of comorbidities (ii) their severity (e.g. for determining the patient's ASA grade) (iii) the type of surgical procedures and (iv) their surgical risk grades, all of this before being able to refer to

Grade 1 Surgery - Adults with ASA 2 with comorbidity from Renal Disease

	age in years			
Test	16 to 40	40 to 60	60 to 80	over 80
Chest X-ray	NO	NO	NO	
ECG	NO			YES
Full Blood Count				
Haemostasis	NO	NO	NO	NO
Renal Function	YES	YES	YES	YES
Random Glucose	NO	NO	NO	NO
Urine Analysis				
Blood gases	NO	NO	NO	NO
Lung function	NO	NO	NO	NO

Fig. 2. Adapted from the NICE preoperative guidelines: investigations are recommended based on: patient (i) age, (ii) ASA, (iii) comorbidities (iv) type of surgical procedure and (v) risk grade of surgical procedure. There are 3 types of result for each test: "test recommended", "test not recommended" and "consider test".

the correct table. All of these tasks are obviously highly knowledge intensive as well as being intellectually demanding. In addition, preoperative health assessors typically see dozens of different patients a day, each with a wide variety of health conditions and scheduled for various types of surgical procedures. In practice, the consequences are that, if in doubt, preoperative investigations would probably be requested regardless of the guideline (i.e. *better safe than sorry*) , thus defeating the purpose of the guidelines in managing efficiently the allocation of preoperative investigation in the health services.

4.2 Investigation Rule Axioms Generation

We combined the use of an ontology and reasoner in the preoperative decision support system in order to automatically make recommendations regarding the suitability of tests based on the NICE guidelines. The first step consisted into transforming the NICE tables into rules. This enabled to considerably reduce overlap and redundant information in the current format of the guidelines. See Table 1 as an illustration of the NICE guidelines as rules. Figure 3 illustrates the same rule seen through the Protégé-OWL development tool [10]. Thus, the 1242 different possible cases currently covered by the NICE guidelines were reduced to around a hundred rules of the type illustrated in Table 1.

As the amount of redundant information in the guidelines is substantial, there are a number of different ways the rule axioms can be expressed: according to comorbidity, surgery grade, ASA grade, etc. Thus we have use a number of heuristics to develop the rules:

Table 1. NICE guidelines as a rule: refers to the first line (Chest X-ray) of the table in Figure 2

Rule Id:	NiceGuidelineRule_ChestXRay_NOT-RECOMMENDED05_N0007
Class name:	ChestXRay_NOT_RECOMMENDED_NT+SurgeryGrade1+ ASA2+Renal_Comorb+Between16_and_80Years
Location	Subclass of: ChestXRay_NOT_RECOMMENDED_NICE_Test
Meaning:	*"Adults between the age of 16 years to 60 years old* *of ASA 2, with renal comorbidities* *undergoing surgery of risk grade 1* *should NOT undergo a chest X-ray investigation* *unless undergoing cardiovascular surgery"*

OWL Axiom	PatientClassificationEntity
	and ChestXRayClassificationEntity
	and (has some (Age
	   and (hasTemporalUnit some Year)
	   and (hasIntNumericValue some int[≥"16"integer,<"60"integer])))
	and (hasASA-gradeValue value "2"integer)
	and (hasAssociatedComorbidity some RenalComorbidity)
	and (hasPlannedSurgicalProcedure some (SurgicalProcedure
	   and (hasSurgicalGrade some SurgicalGrade1)))
	and not (hasPlannedSurgicalProcedure some CardioVascularSurgery)

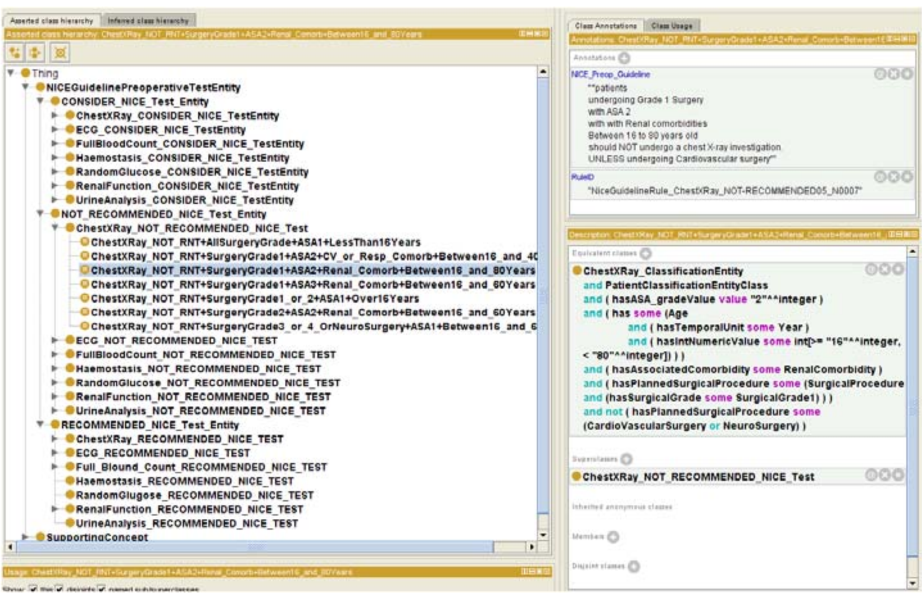

Fig. 3. The NICE Guidelines as OWL Rules as viewed through the Protégé-OWL User Interface

General rules: where possible, we have tried to use as generic rules as possible (e.g. "*All* Adults aged over 80 years old should undergo an ECGTest *regardless* of any other factorsĚ (ASA or surgery gradeĚ))"

Consistency in Rule Generation: Rules can often be generated according to different combinations of factors. For example, the same guidelines could be based on different combinations of factors based on patient ASA, or surgery grade, age, etc. We generally chose to express rules along an interpretable factor whenever possible, such as comorbidities (e.g. *"all patients with a renal comorbidity..."*)

Trade-off between number of rules and complexity of rules: There is a trade-off between an optimum (i.e. minimum) number of rules and their complexity. It is possible to minimise the number of rules by including conjunction, disjunction and exclusion clauses. However, ones needs to ensure that the rules remain interpretable and as close as possible to the original guidelines format. This is mainly for explanation purposes as the decision support tool should be able to provide an understandable explanation of results to end users. Practicaly, this involved chosing to split complex rules into simpler rules when too many factors in the rule made it difficult to interpret. Typically, this involved limiting the number and occurrences of disjunction and exclusion clauses.

The main advantage of modelling the preoperative investigation guidelines as OWL axioms is that the preoperative decision support system can now (i) use third party clinical taxonomies in order to allocate a surgical risk grade to a specific surgical procedure and (ii) use a third party clinical ontology to infer patient comorbidities. Thus, using the OWL patient medical history profile generated at the step 3.b in Figure 1, we can now automatically infer which investigations a patient should have based on his specific medical history.

Figure 4 provides an example of preoperative test recommendation based on reasoning on the decision support ontology. Mark is a 67 year old patient, with arrhythmia, his ASA status has been estimated to be 3 and he is to undergo an open excision of lesion of duodenum. Reasoning on the recommended tests decision support ontology returns that a chest X-Ray, ECG test, full blood count and a renal function tests are all recommended, a haemostasis test may be considered and a lung function test is not recommended. The recommendations are made by the system based on the following reasons:

chest X-Ray: patient is over 60 and of ASA 3, in addition he has arrhythmia which is classified in the decision support ontology as a cardiovascular comorbidity and he is to undergo an open excision of lesion of duodenum, a surgical procedure of grade 4 in the ontology. All this criteria mean that this patient falls within one of 2 categories of patients which are recommended for a chest X-ray investigation.

ECG: test recommended as the patient has a cardiovascular comorbidity so he should undergo an ECG test regardless of all other factors.

Full Blood Count: test recommended as the patient is over 16 years old and is undergoing either grade 3 or 4 surgery.

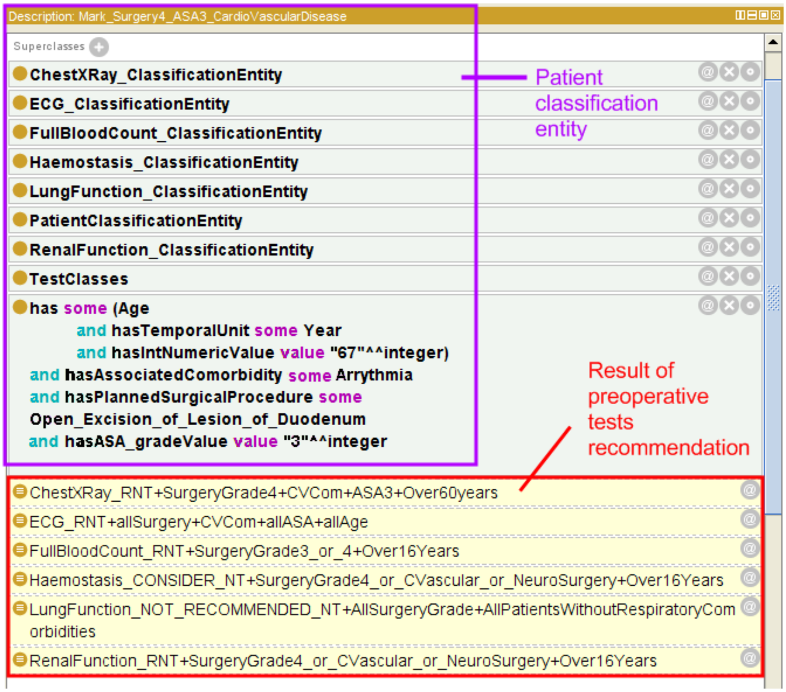

Fig. 4. Example of preoperative test recommendation based on reasoning on the decision support ontology

Renal function: test recommended as patient is over 16 and undergoing grade 4 surgery

Haemostasis: test may be considered as patient is over 16 and undergoing grade 4 surgery

Lung Function: this test is not recommended as the patient does not have any respiratory comorbidities

4.3 Dealing with Multiple Comorbidities

The preoperative guidelines do not explicitly deal with the issue of multiple comorbidities and this is an other area where the decision support tool can provide additional functionalities.

Duplication of test recommendations: In the case of a patient with multiple comorbidities, it is possible that a test may be recommended for multiple reasons. As an example, a patient of ASA 2, over 60 years old, undergoing grade 2 surgery could be recommended an ECG test twice if he has renal comorbidities and cardiovascular comorbidities. In this case, the system can issue a *strong* recommendation alongside relevant explanation.

Conflicting test recommendations: the guidelines are not mutually exclusive and especially not in the case of multiple comorbidities. It is possible for example for a patient of ASA 2, less than 40 years old, undergoing grade 2 surgery to be recommended an ECG test if he has cardiovascular comorbidities but not if he has respiratory comorbidities. The contradiction is only apparent: what is not necessary for a patient with *only* respiratory comorbidities obviously becomes necessary if the patient *also* has cardiovascular comorbidities. Thus, one instance of a "recommended test" within a batch of test results lead to a positive test recommendation regardless of all other test recommendations. According to the same principle, if the system returns instances of "consider test" along "test not recommended" instances, then the system issues a "consider test" recommendation. Finally, the system issues a "test not recommended" advice only if all instances returned are negative for this specific test.

5 Discussion

The work on developing the GuideLine Interchange Format by Ohno-Machado et al. [6] contains a very interesting discussion on the motivation behind the development of GLIF and the process and issues encountered while translating paper-based guidelines into computer-based guidelines. The authors emphasise that this transcription process will usually reveal inconsistencies or flaws in the guidelines. We agree that this is essentially a virtuous cycle as inconsistencies in the guidelines may be highlighted and addressed. Although we found a very small number of inconsistencies in the preoperative guidelines, this did not prove much of a serious issue in our case. A recurring remark about the preoperative guidelines is that they are not exhaustive: they do not cover all possible combinations of surgery, comorbidities and ASA cases. Our understanding is that the guidelines cover a large number and possibly the majority of cases (although we can not confirm this at the time of writing) of patients presenting for elective surgery. Most serious cases are not covered by the guidelines, but this is likely to be when they would be the least helpful, as in these situations, health professionals will use their own clinical judgement to request all relevant preoperative tests or even perhaps seek an alternative to surgery. Our experience would suggest that the major obstacle to effective use of the guidelines in the health services in the format in which they are represented, as they remain both intellectually demanding and knowledge-intensive. This is where a computer-based decision support system can make a genuine difference in practice.

Getting the right format is without doubt a very difficult challenge for the issuers of the guidelines. To be useful, the guidelines need to be comprehensive in covering a large number of cases, as well as being systematic in the presentation of the results, while guaranteeing a rigorous path to a given recommendation for the safety of the patients. Somehow, the format may then take precedence over the underlying *meaning* behind the guidelines. Effectively, since the preoperative guidelines use 5 types of information items as input (surgery grade, surgery type, ASA grade, type of comorbidities and age), these factors are always present in the formulation of the guidelines, even if they are not necessarily relevant in every context. As an example: if a patient has a renal comorbidity, the guidelines recommend a renal function test regardless of any other factors (age, other comorbidities, surgery, etc.) However, this information is not necessarily clear, as it is both repeated and dispersed across multiple tables. Thus, the current format of the guidelines does not convey this information in an efficient manner. An OWL representation of the guidelines make these relations explicit, and is thus perhaps a closer representation of the *intended* meaning behind the guidelines.

6 Conclusion and Future Work

We discussed our experience of using a representation of a patient's medical history in OWL, combined with a reasoning tool to suggest appropriate preoperative tests based on an implementation of preoperative assessment guidelines. We illustrated the reasoning functionalities of the system with a number of practical examples. The system described can reuse third parties ontologies and taxonomies in order to provide advanced decision support functionalities. Future work will include deploying the system in selected pilot sites and observing and analysing how health professionals use the recommendations made by the system. Of particular interest will be to see if health professionals show an interest in the system underlying knowledge representation models and whether these relate to their personal understanding of the guidelines. In other words, we will observe whether the knowledge representation formalisms promote the use and understanding of the guidelines or whether they are of little interest to the health professionals who may be more concerned with the applicability and reliability of the results of the decision support system.

References

1. Ferrando, A., Ivaldi, C., Buttiglieri, A., Pagano, E., Bonetto, C., Arione, R., Scaglione, L., Gelormino, E., Merletti, F., Ciccone, G.: Guidelines for preoperative assessment: impact on clinical practice and costs. International Journal for Quality in Health Care 17(4), 323–329 (2005)
2. Audet, A., Greenfield, S., Field, M.: Medical practice guidelines: current activities and future directions. Annals of Internal Medicine 113, 709–714 (1990)
3. Hunt, D.L., Haynes, R.B., Hanna, S.E., Smith, K.: Effects of computer-based clinical decision support systems on physician performance and patient outcomes. JAMA, Journal of American Medical Association 280(15), 1339–1346 (1998)
4. Amit, G., Neill, A., McDonald, H., Rosas-Arellano, M., Devereaux, P., Beyene, J., Sam, J., Haynes, R.: Effects of computerized clinical decision support systems on practitioner performance and patient outcomes, a systematic review. Journal of the American Medical Association, JAMA 293(10), 1223–1238 (2005)
5. Leong, T.Y., Kaiser, K., Miksch, S.: Free and open source enabling technologies for patient-centric, guideline-based clinical decision support: A survey. IMIA Yearbook of Medical Informatics, Methods of Information in Medicine 46(1), 74–86 (2007)

6. Ohno-Machado, L., Gennari, J.H., Murphy, S.N., Jain, N.L., Tu, S.W., Oliver, D.E., Pattison-Gordon, E., Greenes, R.A., Shortliffe, E.H., Barnett, G.O.: The GuideLine Interchange Format, a model for representing guidelines. Journal of American Medical Informatics Association, JAMIA 5, 357–372 (1998)
7. Johnson, P.D., Tu, S.W., Musen, M.A., Purves, I.: A virtual medical record for guideline-based decision support. In: Proceedings of the 25th Symposium of the American Medical Informatics Association (AMIA), Washington, DC, US, pp. 294–298 (2001)
8. de Clercq, P.A., Blom, J.A., Korsten, H.H., Hasman, A.: Approaches for creating computer-interpretable guidelines that facilitate decision support. Journal of Artificial Intelligence in Medicine 31(1), 1–27 (2004)
9. OWL, Web Ontology Language (2004), http://www.w3.org/2004/OWL
10. Knublauch, H., Fergerson, R.W., Noy, N.F., Musen, M.A.: The Protégé OWL plugin: an open development environment for semantic web applications. In: McIlraith, S.A., Plexousakis, D., van Harmelen, F. (eds.) ISWC 2004. LNCS, vol. 3298, pp. 229–243. Springer, Heidelberg (2004)
11. Horridge, M., Bechhofer, S., Noppens, O.: Igniting the OWL 1.1 touch paper: The OWL API. In: Proceedings of the third International Workshop of OWL Experiences and Directions, OWLED 2007, Innsbruck, Austria (2007)
12. Bouamrane, M.M., Rector, A., Hurrell, M.: Semi-automatic generation of a patient preoperative knowledge-base from a legacy clinical database. In: Proceedings of the 8th International Conference on Ontologies, DataBases, and Applications of Semantics, ODBASE 2009, OTM 2009 Internet Systems. LNCS. Springer, Heidelberg (to appear, 2009)
13. Bouamrane, M.M., Rector, A., Hurrell, M.: Gathering precise patient medical history with an ontology-driven adaptive questionnaire. In: Proceedings of the 21st IEEE International Symposium on Computer-Based Medical Systems, CBMS 2008, Jyväskylä, Finland, pp. 539–541. IEEE Computer Society, Los Alamitos (2008)
14. Bouamrane, M.M., Rector, A., Hurrell, M.: Ontology-driven adaptive medical information collection system. In: An, A., Matwin, S., Raś, Z.W., Ślęzak, D. (eds.) Foundations of Intelligent Systems. LNCS (LNAI), vol. 4994, pp. 574–584. Springer, Heidelberg (2008)
15. Bouamrane, M.M., Rector, A.L., Hurrell, M.: Using ontologies for an intelligent patient modelling, adaptation and management system. In: Meersman, R., Tari, Z. (eds.) OTM 2008, Part II. LNCS, vol. 5332, pp. 1458–1470. Springer, Heidelberg (2008)
16. Bouamrane, M.M., Rector, A., Hurrell, M.: Development of an ontology of preoperative risk assessment for a clinical decision support system. In: Proceedings of the 22nd IEEE International Symposium on Computer-Based Medical Systems, CBMS 2009, Albuquerque, US. IEEE Computer Society, Los Alamitos (to appear, 2009)
17. Copeland, G., Jones, D., Walters, M.: Possum: a scoring system for surgical audit. British Journal of Surgery 78(3), 355–360 (1991)
18. Sirin, E., Parsia, B., Grau, B.C., Kalyanpur, A., Katz, Y.: Pellet: A practical OWL-DL reasoner. Journal of Web Semantic 5(2), 51–53 (2007)