Use of Description Logic Classification to Reason about Consequences of Penetrating Injuries

Daniel L. Rubin, MD, MS, Olivier Dameron, PhD, and Mark A. Musen, MD, PhD
Stanford Medical Informatics, Stanford University School of Medicine, Stanford, CA 94305, USA

Abstract

The consequences of penetrating injuries can be complex, including abnormal blood flow through the injury channel and functional impairment of organs if arteries supplying them have been severed. Determining the consequences of such injuries can be posed as a classification problem, requiring a priori symbolic knowledge of anatomy. We hypothesize that such symbolic knowledge can be modeled using ontologies, and that the reasoning task can be accomplished using knowledge representation in description logics (DL) and automatic classification. We demonstrate the capabilities of automated classification using the Web Ontology Language (OWL) to reason about the consequences of penetrating injuries. We created in OWL a knowledge model of chest and heart anatomy describing the heart structure and the surrounding anatomic compartments, as well as the perfusion of regions of the heart by branches of the coronary arteries. We then used a domain-independent classifier to infer ischemic regions of the heart as well as anatomic spaces containing ectopic blood secondary to the injuries. Our results highlight the advantages of posing reasoning problems as a classification task, and leveraging the automatic classification capabilities of DL to create intelligent applications.

Introduction

There are different approaches to creating medical reasoning systems, such as rule based, probabilistic, and ontological representations. A challenge in creating new decision support systems is to incorporate medical knowledge and to apply that knowledge in flexible ways [1]. Ontologies are a core building block for intelligent applications, being used to represent knowledge and permit knowledge reuse. However, most applications using ontologies for reasoning use only a few specific ontologies. There are few examples in which ontologies have been extensively reused by different reasoning applications or even shared by different components within the same application.

In most reasoning systems that use ontologies, the knowledge used to guide reasoning (control knowledge) is embedded in the application code or in rules used in conjunction with the domain ontology [2]. Control knowledge and domain knowledge are in different places, making application maintenance and development of future extensions cumbersome. In order to extend or create reasoning applications, it is often simplest to create new ontologies or new application software.

We believe that it is advantageous to use description logics (DL; http://dl.kr.org/) in biomedical applications to represent both the domain knowledge and the control knowledge needed for reasoning. Thus, we construe the reasoning problems in the domain as classification tasks. The reasoning application can also leverage high-performance classifiers to reason with the DL. This approach is advantageous because the classifier would not require any domain-specific knowledge, it could be reused in many other applications, and these applications would be easier to maintain.

To test our hypothesis, we chose the problem of reasoning about penetrating injury in the vicinity of the heart. Given a set of anatomic structures that are directly injured by a projectile, we want to create a reasoning application that deduces secondary injuries of two types: (1) regions of myocardium that will be ischemic if a coronary artery is injured, and (2) propagation of injury as bleeding occurs into damaged anatomic compartments that surround the heart.

These reasoning tasks are complex, and it would be beneficial to develop modular reasoning services that address individual components of the larger problem. Thus, we need a knowledge representation and reasoning approach that is amenable to modular development, permits ontology reuse, and avoids embedding reasoning knowledge in the application code.

The Web Ontology Language (OWL) [3] was recently recommended by the World Wide Web Consortium (W3C) as a standard language for the Semantic Web [4]. It is similar to other ontology languages in that it can capture knowledge by representing the concepts and relationships among them. In addition, the OWL-DL[1] version of the language provides support for description logics (DL). Because it is suitable for both knowledge representation and automated reasoning, OWL may be advantageous in creating intelligent applications.

The use of OWL in biomedical applications has to date focused on "terminological" aspects of knowledge; the formal semantics of DL have been used to infer classification taxonomies and to help identify inconsisten-

[1] Hereafter, we will be referring to OWL-DL when using the term "OWL."

cies [5]. We believe that classification as a reasoning method can also be suitable and advantageous in other types of intelligent applications provided the reasoning task can be posed as a classification problem.

We describe our approach to creating reasoning services that fulfill the above desiderata using OWL. In this work we exploit the automated reasoning capability provided by OWL. We explore the capabilities and advantages of using OWL for knowledge representation and reasoning in our application domain.

Methods

Domain knowledge

The key knowledge required for reasoning about the consequences of penetrating injury is anatomic and physiological knowledge. The Foundational Model of Anatomy (FMA) is a frame-based ontology containing the concepts and relationships that pertain to the structural organization of the human body [6]. The Foundational Model of Physiology (FMP) is an evolving symbolic representation of biological functions pertaining to human physiology to support knowledge-based applications that call for physiological knowledge [7]. We extracted classes and *part-of* relationships from these ontologies pertaining to the heart and surrounding tissue compartments using the Protégé API.

We manually translated the extracted ontologies from FMA [8] and FMP into OWL using the OWL plug-in to Protégé [9]. Classes in the original ontologies had corresponding OWL classes in the OWL ontology. Relationships in the ontologies became OWL properties. Classes that were implicitly disjoint in the FMA were declared disjoint in the OWL model.

In addition to this simple translation, we added additional knowledge that was implicit in the FMA and FMP ontologies, such as declaring the parts of the heart to be merologically (part-of relationships) and taxonomically disjoint. We called the resulting OWL ontology our "base OWL ontology," because we subsequently extended it to implement our two reasoning services.

We extended the base OWL ontology in two ways to create two different reasoning services: (1) infer regions of heart damage secondary to coronary artery injuries, and (2) infer the propagation of initial injury caused by bleeding into breached anatomic compartments. The base OWL ontology was extended to create these applications by adding class restrictions and defined classes to represent additional anatomic and physiological knowledge needed by our application but not available in the FMA and FMP. Thus, we reused the original knowledge representation of anatomy and physiology, and we developed two different reasoning services in a modular manner.

Reasoning about coronary artery ischemia

We created a reasoning service to infer the myocardial ischemic consequences of coronary artery injury ("Cardiac Ischemia Reasoner"). We added necessary and sufficient conditions to classes in our base OWL ontology of anatomy to encode the dependency of downstream arterial branches on the upstream arteries, and to represent the regions of the heart myocardium supplied by the coronary artery branches. For example, to represent the coronary arteries that supply the lateral part of the wall of the left ventricle, we add restrictions to the class LateralPartOfWallOfLeftVentricle that specify values for the isSuppliedBy property, such as LeftCircumflexArtery (Figure 1).

An organ may be supplied by more than one artery, in which case damage to one of the feeding arteries will

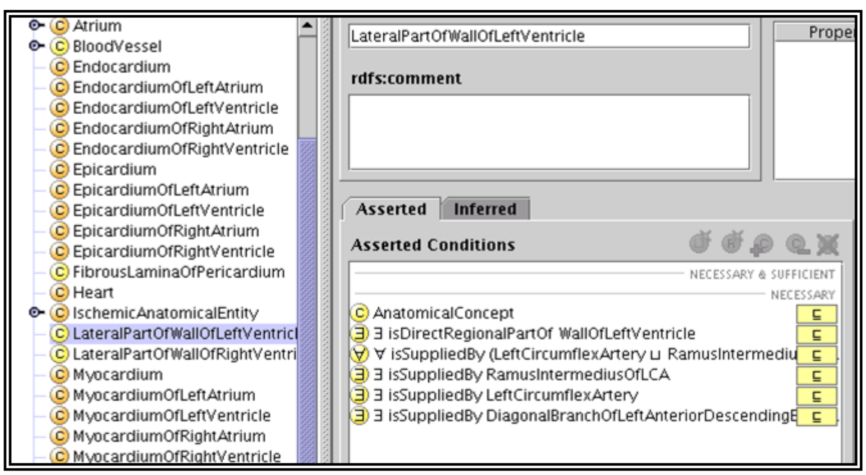

Figure 1. OWL Ontology of coronary anatomy and regional myocardial perfusion. *Classes of anatomic structures are shown in the left panel, and logical definitions of the concepts are on the right. The class* LateralPartOfWallOfLeftVentricle *contains six restrictions representing the necessary conditions for this class. Some of these assertions specify the coronary arterial branches that supply this structure.*

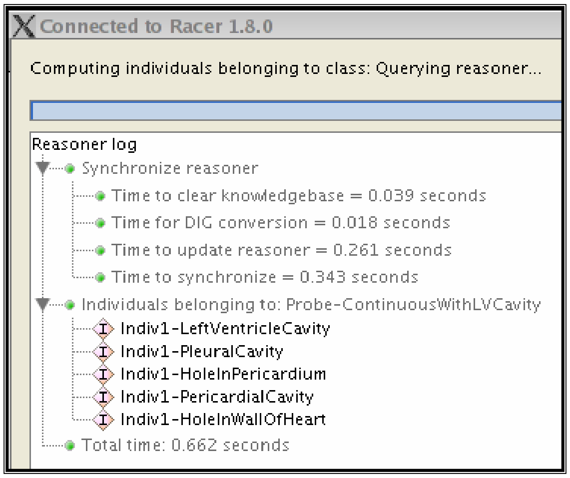

Figure 3. Inferred knowledge after asserting a cardiac injury comprising a hole in the left ventricle and classifying the Injury Propagation OWL ontology. *The pericardial cavity and pleural cavity are inferred to be in continuity with the left ventricle.*

cause partial (not complete) impairment of blood flow to the organ. To represent these types of ischemia, we defined the classes IschemicAnatomicalEntityPartially and IschemicAnatomicalEntityTotally.

This representation in OWL permits automatic reasoning about myocardial regions that become ischemic if an arterial branch supplying part of the myocardium is impaired. To assert an arterial injury from penetrating trauma, one would make the class corresponding to the injured artery a subclass of SeveredBloodVessel. We can then use Racer [10], a domain-independent classifier, to infer a new taxonomy derived from the logical definitions in the original asserted ontology, which includes knowledge about the arterial injury. After classification, results of inference can be read from the ontology by looking for subclasses of the classes that represent the types of injury of interest.

<u>*Reasoning about injury propagation*</u>

We created a second reasoning service to deduce propagation of injury—the development of abnormal conduits that connect anatomic compartments and collections of blood in the pericardium and pleura ("Injury Propagation Reasoner").

To create this reasoning service, we extended the base OWL ontology with knowledge about the anatomic compartments that normally contain blood (heart chambers) and those that do not (pericardial and pleural space), as well as the circumstances under which blood will flow between anatomic compartments. For example, to represent the knowledge that the cavities of the atria and ventricles are filled with blood, we add the following assertion to the class representing the left atrium: LeftAtriumCavity ⊑ [AtriumCavity ⊓ (∃ filledWith.Blood)]. Additional restrictions were added to define particular types of propagated injuries, such as hemopericardium and hemothorax (defined as accumulation of blood in the pericardial space and pleural sac, respectively).

In order to represent a perforation in the wall of the heart wall, we create an instance of the class AddedConduit, and add values to the continuousWith property to describe that this conduit is continuous with both the cavity of the left ventricle and the pericardial space (Figure 2). The continuousWith property represents continuity between adjacent anatomic structures that have been injured, and it is symmetric and transitive. These two property characteristics are needed to infer that, given a perforation in the wall of the left ventricle ("HoleInWallOfHeart") and pericardium ("HoleInPericardium") creating conduits that connect them, the conduits, pericardial cavity, and pleural cavity will be in continuity with the cavity of the left ventricle (Figure 3).

Conduits connecting adjacent anatomic compartments and producing continuities between them will lead to abnormal accumulation of blood in those compartments. Our OWL ontology models the fact that bleeding occurs into anatomical compartments that are connected to other compartments filled with blood. Thus, given a conduit in the wall of the left ventricle and pericardium, our OWL ontology will infer that the pericardial cavity and the pleural cavity will contain blood, in addition to the left ventricle itself. To model pathological states, our ontology contains the class AnatomicalConceptWithEctopicBlood, representing any structure abnormally filled with blood. To define this concept, we created a property filledWithAbnormally, a sub-property of filledWith, representing abnormal blood accumulation.

We performed a preliminary qualitative evaluation of the the capabilities of the Cardiac Ischemia Reasoner and the Injury Propagation Reasoner by creating several coronary artery injury scenarios and presenting the

Figure 2. Knowledge representation in OWL of a hole in the heart wall. *An instance of the* AddedConduit *class is created, having values of the* continuousWith *property specifying the anatomic compartments that this conduit connects.*

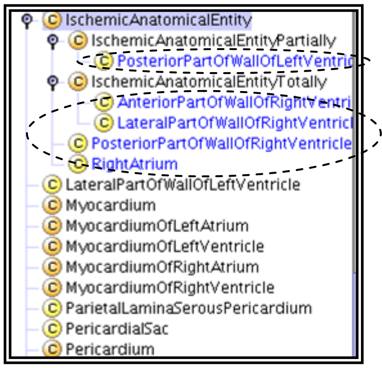

Figure 5. Cardiac Ischemia OWL ontology updated with the knowledge that the second segment of the right coronary artery has been injured. *After automatic classification, particular anatomic classes (circled) are reclassified, suggesting the ischemic regions of myocardium that occur as a consequence of the right coronary artery injury.*

results of automated reasoning to a physician to review in terms of credibility.

Results

In order to use our OWL ontologies to deduce the consequences of penetrating injuries, we create classes and instances in the ontology that describe the new state of knowledge given the injury. For the Cardiac Ischemia Reasoner, we describe an injury such as damage to the second segment of the right coronary artery by making the class SecondSegmentOfRCA a subclass of SeveredBloodVessel. After a classifier is applied to the OWL ontology to re-classify the concepts given this new knowledge, we can deduce which arterial segments and myocardial regions downstream from the primary injury will be damaged by looking for subclasses of IschemicAnatomicalEntity (Figure 5). For example, if the second segment of the right coronary artery is injured, then classes representing myocardial regions that would become ischemic would be classified as subclasses of IschemicAnatomicEntityTotally and IschemicAnatomicEntityPartially (Figure 5). These classes represent complete and partial ischemia, respectively, and depend on the existence of collateral blood supply to myocardial regions. In this manner, the classifier and OWL ontology indicate the regions of myocardium that would be ischemic secondary to the injury.

For the Injury Propagation Reasoner, we create instances that describe the injury and re-classify the OWL ontology to infer the propagated injuries. For example, if the injury is a complete perforation of the left ventricle wall, we update the OWL ontology with this new knowledge by creating an instance of AddedConduit, and declare that this conduit is continuous with both the cavity of the left ventricle and the pericardial space (Figure 2). If the injury also perforates the pericardium, we create an additional instance of AddedConduit for the pericardial injury. After reclassifying the ontology, we can infer the propagated injuries expected in the injured person by looking for instances of the AnatomicalConceptWithEctopicBlood class; in this case, the reasoner deduces that there will be ectopic blood in the pleural cavity and pericardial cavity (Figure 4).

Both of our OWL reasoning ontologies were derived from the same base OWL ontology. Thus, we were able to share the same knowledge model among two different reasoning applications. In addition, while creating the reasoning applications, we found it helpful to have both the domain knowledge and classification knowledge (class restrictions) represented declaratively in the same ontology as we extended and refined our models.

We have successfully applied this reasoning service to infer the effects of coronary-artery injuries given a variety of coronary artery segments as well as to predict the propagated injuries occurring secondary to perforating and non-perforating heart injury, confirmed by a physician who reviewed the results.

Discussion

The choice of representation formalism is a trade-off between the expressivity required for the modeling

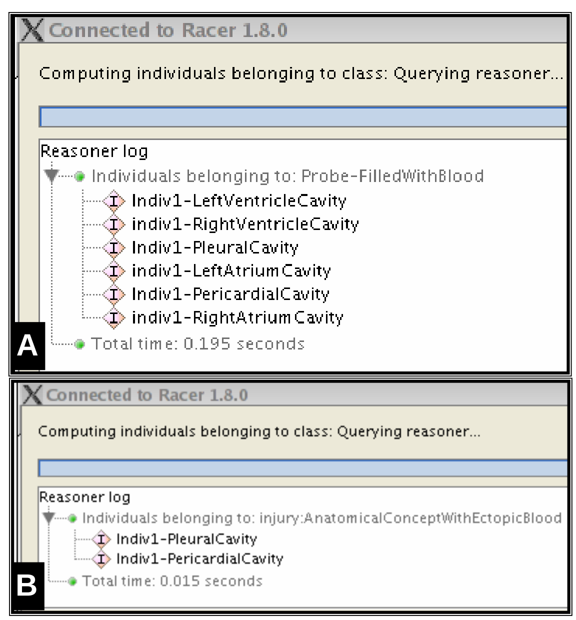

Figure 4. Inference using the Injury Propagation Reasoner.
A) After asserting a conduit in the wall of the left ventricle, the reasoner infers that the pericardial cavity and the pleural cavity are filled with blood, in addition to the cardiac chambers.
B) A defined class is used to infer the propagated injuries—the pericardial cavity and the pleural cavity of the wounded individual contain ectopic blood.

step, and the computational constraints imposed by the reasoning application. Other applications have been developed to perform anatomical reasoning, such as the TraumAid system [11]. They take a constructive approach to problem-solving, and tend to be computationally inefficient and cumbersome to build and maintain.

Our results suggest that inferring the consequences of penetrating injury can be formalized as a classification task. There are benefits in using OWL as a representation language. OWL is an emerging standard, and reasoning applications can take advantage of high-performance classifiers such as Racer. OWL ontologies contain both a declarative model of the domain knowledge as well as explicit class definitions, properties, and axioms that specify the knowledge used in the classification task. Since all knowledge needed for reasoning is in the ontology, the application code can be reused among different reasoning tasks without modification. In addition, we were able to model our reasoning tasks in OWL simply by adding a few new classes and axioms to the base OWL ontology—we did not need to develop specialized reasoning tools.

In creating our two reasoning applications, we translated a subset of the FMA (stored as a frame representation in Protégé) into OWL. The initial part of this translation is simply a matter of changing syntax of the knowledge representation; however, there is also a change in semantics because of the open-world assumption in OWL. Thus, we also needed to add additional class restrictions such as closure axioms, to deal with the change in semantics.

We found that separating the reasoning problem into two discrete tasks was helpful, because we could divide the larger problem into smaller modular components and work on them separately. In the future, we could combine the two OWL ontologies into a single ontology, particularly since both ontologies share the same base ontology. This modular approach to OWL ontology representation and reasoning service development could also be useful in tackling larger reasoning applications.

Our current reasoning applications can infer cardiac ischemic damage and propagated injuries given a list of damaged anatomic structures (Figure 5 and Figure 4). This reasoning is limited since it focuses only on the region of the heart. However, our work demonstrates the principles of an approach that we believe can be generalized to the remainder of the torso. Another limitation of our methodology is that reasoning is deterministic, yet predicting injuries is clearly a task fraught with uncertainties. Symbolic reasoning approaches such as description logics constrain us to deterministic classification results.

In conclusion, we have demonstrated benefits of OWL as a representation language for reasoning in an intelligent biomedical application. It may be helpful in other reasoning applications and could improve the ability to share and reuse domain and reasoning knowledge among reasoning applications.

Acknowledgments

This work was supported by a contract from DARPA, executed by the U.S. Army Medical Research and Materiel Command/TATRC Cooperative Agreement, Contract W81XWH-04-2-0012. This work was also supported by the Protégé resource, under grant LM007885 from the U.S. National Library of Medicine.

References

[1] Haimowitz IJ, Patil RS, Szolovits P. Representing medical knowledge in a terminological language is difficult. In: Proceedings of the Symposium on Computer Applications in Medical Care; 1988: IEEE Computer Society Press; 1988. p. 101–105.
[2] Golbreich C. Combining rule and ontology reasoners for the Semantic Web. Rules and Rule Markup Languages for the Semantic Web, Proceedings 2004;3323:6-22.
[3] Smith MK, Welty C, McGuinness DL. OWL Web Ontology Language Guide. http://www.w3.org/TR/owl-guide/ 2003.
[4] WorldWideWeb Consortium. OWLWeb Ontology Language Reference. In; W3C Recommendation 10 Feb, 2004.
[5] De Coronado S, Haber MW, Sioutos N, Tuttle MS, Wright LW. NCI Thesaurus: Using Science-based Terminology to Integrate Cancer Research Results. Medinfo 2004;2004:33-7.
[6] Rosse C, Mejino JL, Jr. A reference ontology for biomedical informatics: the Foundational Model of Anatomy. J Biomed Inform 2003;36(6):478-500.
[7] Cook DL, Mejino JL, Rosse C. Evolution of a Foundational Model of Physiology: Symbolic Representation for Functional Bioinformatics. Medinfo 2004;2004:336-340.
[8] Dameron O, Rubin DL, Musen MA. Converting the Foundational Model of Anatomy into OWL. In: Proc AMIA Symp; 2005; Washington, DC; 2005 (accpeted).
[9] Knublauch H, Fergerson RW, Noy NF, Musen MA. The Protege OWL Plugin: An open development environment for Semantic Web applications. Semantic Web - Iswc 2004, Proceedings 2004;3298:229-243.
[10] Haarslev R, Möller V. Description of the RACER System and its applications. In: Proceedings of International Workshop on Description Logics (DL-2001); 2001; Stanford, CA USA; 2001. p. 1-3.
[11] Clarke JR, Hayward CZ, Santora TA, Wagner DK, Webber BL. Computer-generated trauma management plans: comparison with actual care. World J Surg 2002;26(5):536-8.

Address for correspondence

Daniel L. Rubin (rubin@smi.stanford.edu), Stanford Medical Informatics, MSOB X-215, Stanford, CA 94305, USA.