Artificial Intelligence in Medicine (2006) **37**, 167—176

ARTIFICIAL INTELLIGENCE IN MEDICINE

http://www.intl.elsevierhealth.com/journals/aiim

Using ontologies linked with geometric models to reason about penetrating injuries

Daniel L. Rubin*, Olivier Dameron, Yasser Bashir, David Grossman, Parvati Dev, Mark A. Musen

Stanford Medical Informatics, MSOB X-215, Stanford University, Stanford, CA 94305, United States

Received 15 April 2005; received in revised form 22 March 2006; accepted 23 March 2006

KEYWORDS
Anatomy ontologies;
Computer reasoning;
Geometric heart models

Summary

Objective: Medical assessment of penetrating injuries is a difficult and knowledge-intensive task, and rapid determination of the extent of internal injuries is vital for triage and for determining the appropriate treatment. Physical examination and computed tomographic (CT) imaging data must be combined with detailed anatomic, physiologic, and biomechanical knowledge to assess the injured subject. We are developing a methodology to automate reasoning about penetrating injuries using canonical knowledge combined with specific subject image data.
Methods and material: In our approach, we build a three-dimensional geometric model of a subject from segmented images. We link regions in this model to entities in two knowledge sources: (1) a comprehensive ontology of anatomy containing organ identities, adjacencies, and other information useful for anatomic reasoning and (2) an ontology of regional perfusion containing formal definitions of arterial anatomy and corresponding regions of perfusion. We created computer reasoning services ("problem solvers") that use the ontologies to evaluate the geometric model of the subject and deduce the consequences of penetrating injuries.
Results: We developed and tested our methods using data from the Visible Human. Our problem solvers can determine the organs that are injured given particular trajectories of projectiles, whether vital structures — such as a coronary artery — are injured, and they can predict the propagation of injury ensuing after vital structures are injured.
Conclusion: We have demonstrated the capability of using ontologies with medical images to support computer reasoning about injury based on those images. Our methodology demonstrates an approach to creating intelligent computer applications that reason with image data, and it may have value in helping practitioners in the assessment of penetrating injury.

* Corresponding author. Tel.: +1 650 725 5693; fax: +1 650 725 7944.
E-mail address: rubin@smi.stanford.edu (D.L. Rubin).

0933-3657/$ — see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.artmed.2006.03.006

1. Introduction

Rapid and effective medical intervention in response to civil and military-related injuries is crucial for saving lives and limiting disability [1]. Assessing penetrating injuries is a complex task, requiring both geometric data and anatomic knowledge. Accurate assessment of these injures is challenging because the spatial relationships among anatomic regions can be complex, and potential damage to certain vital structures may not be recognized. Intelligent tools that can integrate patient-specific geometric data and anatomic knowledge to inform care providers about internal injuries can be valuable to patient care.

Geometric data are specific to the patient, usually obtained from volumetric CT images, and they contain spatial information about the size and location of visible structures in the body. Anatomic knowledge adds meaning and insight to geometric data, labeling regions in space with particular organs, relating organ parts and subparts to other anatomic structures, and identifying critical structures that may affect patient prognosis and management.

Geometric data alone are insufficient to assess penetrating injury. The resolution of imaging modalities is limited, and small anatomic structures may not be visible in the images. Certain organs are known to be adjacent to others, but this knowledge is not contained in the images—an expert interpreter of the images is needed. Injuries to some regions of anatomy (such as an artery supplying the heart) will result in damage to other organs that were not directly injured by a projectile (the wall of the heart that is supplied by the injured artery). Such knowledge is not contained in the geometric data, but is known to domain experts.

A project to use both geometric data derived from images and canonical anatomic knowledge to aid the rapid diagnosis of penetrating injury called the Virtual Soldier project is being undertaken by the U.S. Defense Advanced Research Projects agency [2]. The vision for the project is that each soldier could wear a ''dog tag'' containing both pre-injury CT images and other relevant baseline clinical data. At the time of an injury, an information system would read the images off the dog tag and offer advice about the nature of the wound, the patient's prognosis, and requirements for therapy.

In order to fulfill the vision of this project, we need to develop a method to integrate patient-specific geometric data and anatomic knowledge and use that knowledge to reason about penetrating injuries. In this paper, we describe our approach to this integration problem and our development of problem solving services to use anatomic knowledge to reason qualitatively about penetrating injuries.

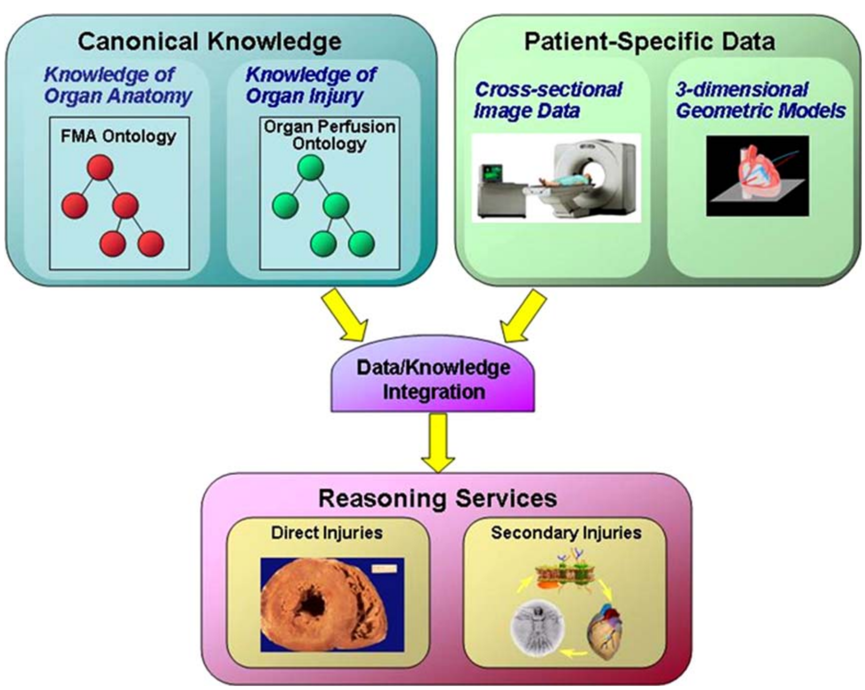

Figure 1 Architecture for integrating patient-specific data and canonical knowledge to reason about penetrating injury.

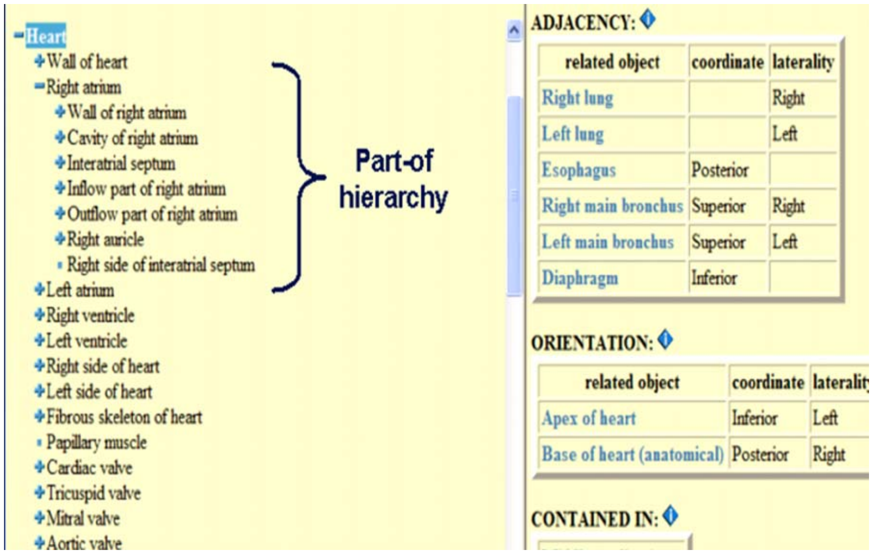

Figure 2 The Foundational Model of Anatomy (FMA). A portion of the FMA ontology is displayed using a Web interface, showing structures related to the heart. Organs and organ parts are shown in the hierarchy on the left (a ''partonomy'' display). Knowledge about individual organs or organ parts is shown in the panel on the right, and includes information such as adjacencies, orientation, and containment. The ontology is stored in Protégé.

2. Methods

The architecture of our system to integrate patient-specific geometric data with anatomic knowledge in ontologies is shown in Fig. 1. Our approach is to perform qualitative reasoning using canonical anatomic knowledge in ontologies combined with patient-specific geometry to infer the identity of injured organs and likely consequences of the organ injuries. Canonical knowledge sources contain detailed knowledge of organ anatomy as well as knowledge about structural anatomic dependencies that are important for predicting secondary injuries. Patient-specific data comprise cross-sectional imaging data and three-dimensional geometric models that are built from these data. Data structures in our software architecture integrate the canonical knowledge and patient-specific geometric data, making both available to applications (reasoning services) that can perform intelligent tasks such as predicting direct and secondary injuries (Fig. 1).

2.1. Knowledge sources

2.1.1. Ontology of canonical anatomy

Rosse and colleagues [3] have developed a comprehensive ontology of human anatomy known as the Digital Anatomist Foundational Model of Anatomy (FMA). The FMA contains more than 70,000 entities that describe the elements of canonical human morphology in a clear and consistent manner (Fig. 2). The ontology is modeled using the Protégé ontology-management environment (http://protege.stanford.edu[1]), and it adheres to the conventions of the OKBC frame language [4]. With this representation, reasoning services access knowledge in the FMA by locating the pertinent anatomic entities (''classes'' in the ontology), and reading their attributes (''slots'' on the class). Slots can be atomic types (such as integers, strings, etc.) or other classes (e.g., the ''part-of'' slot contains classes that have a partonomic relationship with the given class).

The FMA provides declarative descriptions of detailed anatomic structures in a computationally accessible format. Knowledge in the FMA that we use in this project includes organ names, compositionality (partonomy relationships), organ adjacencies, containment, and continuities (Fig. 2).

The FMA is particularly useful because it contains anatomic structures that may be too small to be visible on images, and thus may not be present in geometric models. This knowledge is useful for a reasoning service to deduce possible injury to small structures that are adjacent to visible structures.

2.1.2. Ontology of coronary anatomy and regional perfusion

While the FMA is an excellent knowledge source describing morphology and composition of anatomic structures, it lacks physiological and pathophysiological knowledge. In particular, it does not describe the regions of myocardium supplied by branches of

[1] Accessed 21 March 2006.

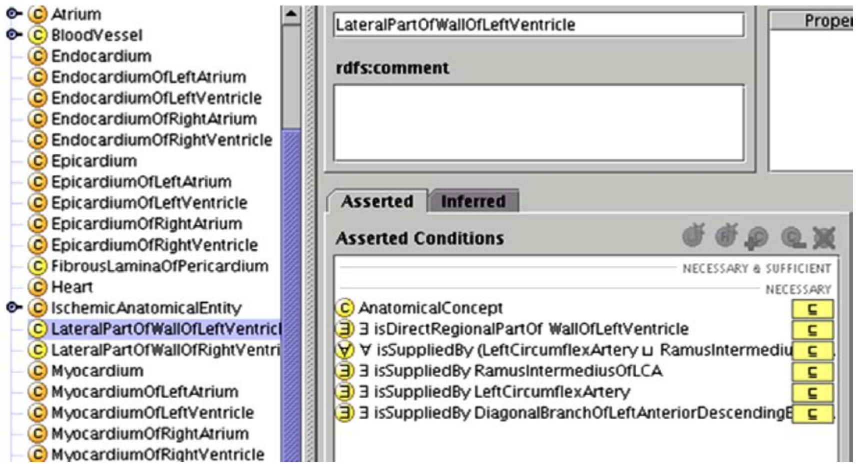

Figure 3 Ontology (in OWL) of coronary anatomy and regional myocardial perfusion. Classes of anatomic structures are shown on the left panel, and formal definitions of the entities are shown on the right. The class ''LateralPartOfWallOfLeftVentricle'' is seen to be defined by six assertions, all necessary conditions for this class. Some of these assertions specify the coronary arterial branches that supply this structure.

the coronary arteries. Such knowledge is needed to reason about secondary organ damage—injury that occurs to particular anatomic structures as a result of damage to other structures.

We built an ontology of coronary artery anatomy and regional myocardial perfusion (Fig. 3) using the Web Ontology Language (OWL) [5]. The OWL classes contain formal definitions, represented using logical statements, specifying the necessary and sufficient conditions (''assertions'') for class inclusion. For example, the definition of the lateral wall of the left ventricle includes assertions specifying all of the branches of the coronary arteries that ordinarily supply it (Fig. 3). This ontology specifies the segments and continuities in coronary arteries, the composition of myocardial regions (e.g., the left ventricle has anterior, lateral, posterior, apical, and septal parts), and it describes the myocardial regions supplied by particular coronary arterial branches. This ontology also defines the coronary arteries as being ''critical'' structures—anatomic structures that result in damage to other structures if they are injured.

The class definitions contained in the OWL ontology permit reasoning services to deduce important physiological consequences of arterial injury. First, the ontology encodes the knowledge that arterial branches downstream from an injured branch will be functionally impaired. Second, the ontology contains knowledge of all arterial branches feeding a myocardial region. There are also definitions about when a region is totally or partially ischemic (a region is totally ischemic if all arteries supplying it are impaired, and is partially ischemic if one ore more arteries are not impaired).

Our OWL ontology does not explicitly model time. Physiological consequences of injury evolve over time, and our approach is to model time in terms of states. Our OWL ontology initially contains assertions pertaining to canonical anatomy and physiology (''pre-injury state''). Immediately following an injury, the anatomy is altered, and this knowledge is represented by making new assertions in the OWL ontology (''injury-asserted ontology'') to describe the injury (''post-injury state''). Sometime following the injury, the primary injury will propagate (''injury propagation state''). Both the post-injury and injury propagation consequences of injury are inferred from the OWL ontology by examining which classes are re-classified to new parent classes when the classifier is applied to the injury-asserted ontology.

Protégé provides support for OWL development, so we were able to coordinate the development of the OWL ontology of cardiac perfusion with the FMA in the same ontology development environment.

2.2. Geometric data sources and model

We obtained segmented images from the Visible Human project [6]. These comprise serial cross-sectional images from a cadaver, and they are analogous to reconstructed images available from CT on live patients. The images had been manually segmented, a process in which non-overlapping geographic regions in the raw images are assigned labels identifying anatomic structures. These labels were used to map these anatomic structures to corresponding anatomic classes in our ontologies.

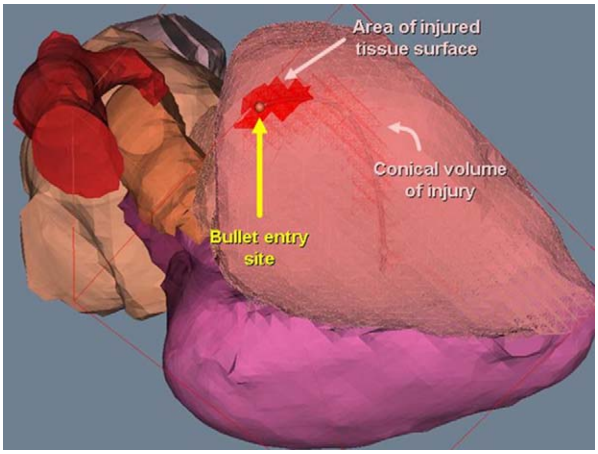

Figure 4 Three-dimensional geometric model of the heart with labeled anatomic structures (shaded volumes in the geometric model correspond to anatomic structure classes in the FMA ontology). A trajectory of penetrating injury is superimposed (curved tubular shaded area). A conically-shaped region of tissue injury is predicted and displayed in the geometrical model (conically-shaped shaded region shown by arrow). We can determine the identity of injured anatomic structures and the volume of damaged organs from this geometric model. We can infer possible injuries to adjacent structures using knowledge in the FMA ontology.

The segmented two-dimensional images cannot be used directly for spatial anatomical reasoning; a three-dimensional representation of patient anatomy must be built from these images to reason with a three-dimensional projectile trajectory. We used the Insight Toolkit (ITK; http://itk.org, see Footnote 1) to build solid three-dimensional tetrahedral mesh models from the serial segmented images of the chest (Fig. 4). These geometric mesh models created from the imaging data represented the three-dimensional coordinates of anatomic structures in space. Collections of vertices in the mesh model were labeled with FMA class names to identify the anatomic structures that they represent. This task is accomplished in ITK by creating data structures that contain collections of vertices representing particular structures (called "spatial objects"). The spatial objects were extended to include the name of the FMA class of the anatomic structure represented by the spatial object. In this manner, the mesh model of patient-specific geometry is linked to canonical anatomic structures in the FMA (as well as the OWL ontology of myocardial perfusion). Our methodology provides a software architecture that makes patient-specific geometric data and canonical anatomic knowledge accessible to intelligent applications such as reasoning services.

We created a graphical visualization application that displays patient-specific geometric data models. Spatial objects comprising sets of tetrahedrons that represent particular organs or organ parts are displayed in different colors. A specified trajectory of penetrating injury can be incorporated into the geometric model as an additional spatial object. Rendering methods are applied to highlight the surface regions and internal volume of organs affected by the penetrating injury and the areas surrounding it (Fig. 4).

2.3. Describing projectile trajectory and wounds

Given a patient-specific geometric model, we can describe and represent any arbitrary trajectory of injury by specifying two coordinates: an entry point and exit point. We assume the following scenario: the injured subject will have had a baseline volumetric scan of the body to capture patient-specific geometry. In our case, the injured "subject" is the Visible Human, since our input geometric data set is from the Visible Human. In practice, the input data set for geometry would be a CT scan of the subject.

We assume that the care giver attending to the injured subject can locate these two points on the body and specify them by relating them to external body landmarks. Once the entry and exit wounds have been specified, we specify the trajectory of the projectile as a linear path connecting entry and exit points. This trajectory is used by the subsequent reasoning services as input to infer direct and secondary injuries.

2.4. Reasoning services

We have initially implemented two applications that use patient-specific geometric data and our canonical knowledge sources to perform useful reasoning capabilities: (1) a tool to determine which organs are injured by a penetrating injury ("Direct Injury Reasoner") and (2) a tool that determines whether any vital structures are injured and the consequences of such injury ("Secondary Injury Reasoner").

The Direct Injury Reasoner takes as input an entry wound and an exit wound on the patient, and it deduces the anatomic structures that have been injured by direct impact by the penetrating injury (or from shock waves in close proximity to the trajectory of injury). To accomplish this task, the Direct Injury Reasoner defines a parametric trajectory path of the penetrating injury using the observed wounds and three-dimensional tetrahedral mesh model of the patient derived from the image data. This trajectory is added to the geometric model, and it is used to infer the region of injury created by a projectile (Fig. 4). The Direct Injury Reasoner deduces the anatomic structures that have been hit by the projectile by identifying the set of geometric elements intercepted by the trajectory and mapping them to the corresponding entities in the FMA. It also creates a parametric representation of a region of tissue damage in the vicinity of the trajectory due to shock waves and tissue strain. The Direct Injury Reasoner can only infer damage to structures that are visible in the segmented images that were used to build the geometric model of the patient.

By linking the geometric mesh model to the FMA, the FMA's rich set of relationships among anatomic entities become accessible to the Direct Injury Reasoner so that it can deduce injury to structures not visible in the segmented images. The FMA includes complex adjacency relationships (noting, for each structure, other structures that may be superior, inferior, anterior, posterior, to the left, or right), orientation relationships (e.g., that the apex of the heart is inferior and to the left), and contained-in relationships (e.g., that the heart is contained in the middle mediastinum). The linkage between the geometric model and the FMA allows the Direct Injury Reasoner to consider the path of a penetrating injury in anatomic terms and to deduce which adjacent anatomic structures not visible in the image data may have also been injured.

The Secondary Injury Reasoner takes as input a list of anatomic structures that have been injured by the penetrating injury (deduced by the Direct Injury Reasoner) and it deduces additional tissue injuries that are occurring or will occur as a consequence of the primary injuries. To accomplish this task, the Secondary Injury Reasoner first examines the injured organs to identify critical anatomic structures. At this point in our work, we have modeled coronary arteries as the only critical structures. The Secondary Injury Reasoner can recognize which of the structures are critical by querying our OWL ontology, because it defines which anatomic structures are critical (Fig. 3).

If any critical structures have been injured, the OWL ontology is updated with this information by creating new assertions (new subclasses indicating the structures that have been injured). For example, if the second segment of the right coronary artery (RCA) is injured, then a new subclass of "functionally impaired blood vessel" would be created by the Secondary Injury Reasoner (Fig. 5).

After the Secondary Injury Reasoner asserts the damaged critical structures, it calls a classification engine that updates the OWL ontology, inferring new classes and relationships given the asserted knowledge and pre-existing class definitions. Finally, the Secondary Injury Reasoner examines the updated OWL ontology to determine if there is new knowledge about injured organs as a consequence of the critical organ injury previously

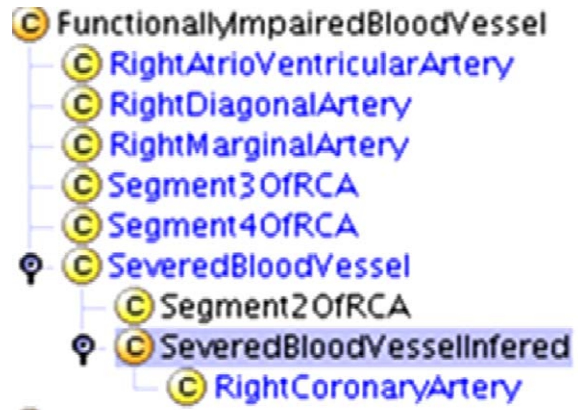

Figure 5 OWL ontology of coronary anatomy and regional myocardial perfusion, with a new assertion that the second segment of the right coronary artery has been injured (a new class, "Segment2OfRCA" was created under "FunctionallyImpairedBloodVessel"). After automatic classification, additional impaired blood vessels are deduced (classes in light color), such as segments 3 and 4 of RCA, which are downstream from the injured segment 2 of RCA.

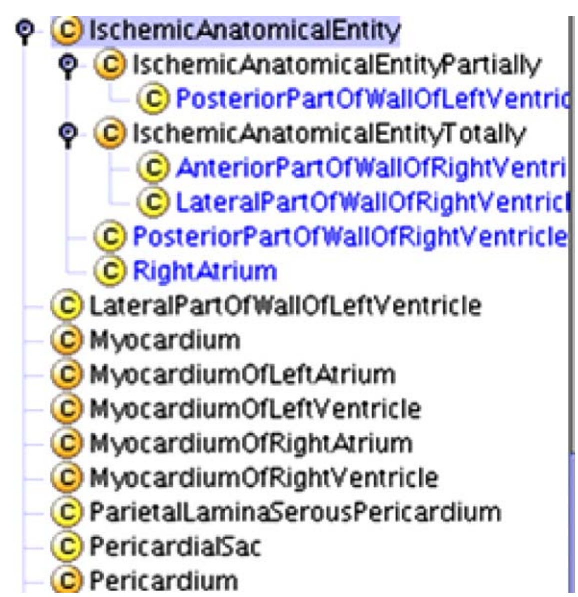

Figure 6 OWL ontology of coronary anatomy and regional myocardial perfusion, updated with the knowledge that the second segment of the right coronary artery has been injured. After automatic classification, new classes (light color) appear as subclasses of "IschemicAnatomicalEntity," suggesting that regions of the left ventricle, right ventricle, and right atrium are ischemic as a result of the right coronary artery injury previously asserted (Fig. 5).

asserted. It accomplishes this task by looking for new subclasses of the OWL ontology that contain classes representing secondary injury (Fig. 6). In our example, the Secondary Injury Reasoner would deduce that most of the right ventricle and portions of the left ventricle and right atrium were ischemic as a result of the injury to the second segment of the right coronary artery.

By combining the Direct and Secondary Injury Reasoners in series, we can begin with baseline patient imaging data and patient wounds and rapidly deduce the anticipated extent of direct and secondary injuries.

3. Discussion

We undertook this work to develop intelligent applications to improve the ability of practitioners to assess and triage injured patients. One might undertake this task simply by creating geometric models of anatomic structures from volumetric images alone. However, this approach would not be adequate to solve this task. Images of patient anatomy may contain labels, but there is little knowledge in those labels beyond a name; thus, it would be impossible to use the image information alone in a computer application to reason about the consequences of penetrating injury. Furthermore, small vital anatomic structures may not be visible on the imaging modality used to create the geometric models. There will be no label in the model for such small structures, and it would not be possible for a reasoning service to deduce those structures they could be injured without access to additional anatomic knowledge.

In our context, *reasoning* means inferring new knowledge from knowledge asserted in the ontologies that are used in the reasoning task. Reasoners are tools that use ontologies and patient data and perform reasoning with it.

In order to support reasoning about the consequences of organ injury, we needed to enhance geometrical data with anatomic knowledge. We needed to go beyond annotating geometric models with organ names and take advantage of the knowledge encoded in ontologies such as the FMA. This approach permits us to develop reasoning services such as the Direct Injury Reasoner that identifies structures adjacent to injured organs that may be injured. It also allows these services to suggest other anatomic structures that may have been injured but that are not visible in the patient images.

Our approach to automated reasoning about the consequences of injuries is a qualitative method based on canonical symbolic representation of knowledge using ontologies. An alternative approach could have been to create quantitative parameterized models that incorporate actual clinical observations such as heart rate and blood pressure. In such quantitative approaches, systems of ordinary differential equations are created to describe the temporal patterns of physiological signals in great detail [7,8]. While such quantitative models were motivated from anatomic models, they have no direct symbolic connection to the anatomic entities that they represent. In our application, we are interested in qualitative physiological reasoning: determining the identity of injured organs and the extent of injury, as well as how injury has propagated. This information is qualitative and is amenable to symbolic representation. For this reasoning task, we need to know the geometry and anatomy, data which we have acquired. Parametric quantitative models would be appropriate if our reasoning task were to determine quantitative physiological changes secondary to the penetrating injuries, or to fit a parametric physiological model to observed clinical parameters. Such reasoning tasks are important, but they are different from the tasks we are undertaking in this study where we know the trajectory and geometry, and wish to perform qualitative causal reasoning about the consequences of those injuries.

We use two different types of ontologies to support the qualitative reasoning services that we developed. The first ontology is the FMA, which encodes a breadth of anatomic knowledge in a declarative format consisting of entity classes and numerous informative relationships and attributes. Ontologies like the FMA specify the entities in the domain and the relationships among them, providing a domain of discourse that is meaningful to both humans and intelligent computer applications.

Some of the tasks of evaluating penetrating injury require additional anatomic and physiological knowledge not contained in the FMA, such as inferring additional anatomic injury secondary to primary injuries. Thus, we developed a second ontology to represent cardiac arterial anatomy and myocardial perfusion, encoded using OWL.

OWL is a language for defining Web ontologies. In OWL, an ontology is a set of definitions of classes and properties, and constraints on the way those classes and properties can be employed [5]. In this representation, we formally specify the meaning of certain anatomic entities and capture their anatomic semantics (we define regions of cardiac myocardium in terms of the coronary artery branches that supply them). Representing this knowledge in OWL makes it accessible to reasoning applications such as automatic classification based on description logics, and the task of reasoning about the consequences of arterial injury is posed as a classification problem. The advantage of this approach is that this reasoning task can be implemented using a domain-independent classifier that is applied to the OWL ontology.

OWL has also been recently recommended by the World Wide Web Consortium (W3C) as a standard language for the Semantic Web [5]. It is similar to other ontology languages in that it can capture knowledge by representing the entities and relationships among them. In addition, OWL 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. We believe that classification as a reasoning method is advantageous in intelligent applications such as ours because the reasoning problem can be posed as a classification problem. Thus, we can take advantage of OWL as an emerging knowledge representation standard and exploit high-performance classifiers for reasoning instead of having to develop a domain-specific reasoning service. Further, using OWL permits us to adapt and extend our knowledge model and refine our reasoning task in conjunction with maintaining and modifying the ontology that contains knowledge of the task, simplifying code development and knowledge management.

Our results suggest that inferring the consequences of penetrating injury can be formalized as a classification task. 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.

Classification is used to infer specialization relationships between classes from their formal definitions. A domain-independent classifier takes an input class hierarchy and the logical expressions it contains, and returns a new class hierarchy, which is logically equivalent to the input hierarchy. In our work, we use this classification approach to infer new knowledge: by asserting a new fact in our OWL representation of coronary anatomy, such as the presence of injury to the second segment of the right coronary artery, we can discover after automatic classification that several regions of the myocardium will be ischemic. Regions of myocardium that are perfused by other arteries that are not occluded are not inferred to be ischemic (Fig. 6). By taking this knowledge modeling approach, we can use the power of automatic classification to cast the problem of reasoning about the consequences of arterial injury as a classification problem.

We used the Protégé suite of tools (http://protege.stanford.edu, see Footnote 1) in this work to manage and access the ontologies. A benefit of using Protégé for ontology management in this project is that it supports both OKBC and OWL representations, and it can invoke automated classification engines. It also provides a Java API for developing applications such as our reasoning services.

A limitation of our current approach is that it does not incorporate uncertainty relating to injuries. Our method is a deterministic symbolic reasoning approach to inferring the consequences of injury. There are several ways that information may be uncertain. If there is more than one possible trajectory, then we cannot know with certainty which trajectory to use for reasoning. One way of handling such situations is to consider the output of reasoning an hypothesis of the state of injuries in the subject, rather than a definite answer. Then in cases of multiple possible trajectories, our reasoning approach could produce a set of alternative hypotheses of injury, and the most likely hypothesis can be selected by evaluating additional clinical information from the subject. We have not yet undertaken such an approach to reasoning using alternative hypotheses, but we believe that this is a strategy for using deterministic symbolic reasoning in the setting of uncertainty.

Another limitation of our study is that we have not yet performed a formal evaluation of our approach. In performing an evaluation, it is important to define when performance is satisfactory. In our case, performance is satisfactory when we correctly identify the organs that are injured. We are now planning to undertake evaluation studies in phantoms in which we are creating internal "organs" of varying size and subjecting the phantom to projectile injury. We will be comparing our reasoning results in those phantoms with observed damage.

While we have shown the utility of computer reasoning about physiological consequences of injury using OWL, we do not claim that all aspects of physiological representation can be achieved using OWL or other DL frameworks. If one desired to have a high-resolution representation of time, then representations other than OWL would probably be preferable. We have represented the temporal dimension of physiological data in OWL in a coarse manner: specifically, we represent time in terms of states: pre-injury, immediately post-injury, and injury propagation at some unspecified time (though usually in close temporal proximity) following the injury. Our approach of binning the temporal dimension into discrete states permitted us to use OWL in a reasoning task that required dealing with the temporal dimension. This choice was appropriate because our reasoning task uses canonical anatomical and physiological knowledge, and we do not use detailed subject-specific knowledge that would be needed to perform high-fidelity temporal physiological reasoning.

Previous work related to assessing penetrating injury has focused on developing simulation environments and teaching aids to assist in assessing penetrating injuries [9]. While such teaching activity is valuable to give practitioners experience managing such trauma cases, it does not replace the need to have specific knowledge about the particular constellation of injuries in the particular patient being cared for.

In other related work, Ogunyemi and colleagues developed a system that calculated the probabilities of organ injuries using a canonical geometric model of human anatomy, and used a Bayesian network to classify particular combinations of injuries [10]. While their approach can be a helpful guide to typical injuries, its geometric models were not specific to the particular patient. In addition, the reasoning capability is a hard-coded classifier. In our work, we adopt ontological representations of knowledge and create reasoning services that use these ontologies and patient-specific geometric models. These reasoning services can be extended or modified without having to re-engineer the underlying knowledge or data representations. For example, Direct Injury Reasoner identifies anatomic structures that intersect the trajectory path of the projectile, and are directly injured. Subsequently, we created the Secondary Injury Reasoner that uses the ontology of myocardial perfusion in OWL to recognize coronary artery segments in the list of injured structures and reason about myocardial ischemic damage that occur as a result of coronary artery injury.

In our approach, the knowledge sources are canonical, while particular patient anatomy can be variable. Our knowledge sources were designed to be "canonical," meaning that they represent the typical entities and relationships that are observed. This limitation can be overcome by extending our ontologies to model anatomic variation. We are currently adding the balanced and right dominant patterns of coronary arterial supply to our myocardial perfusion ontology. In this way, the patient-specific geometric model can be associated with a more appropriate anatomic knowledge source.

In our work, reasoning is focused on the chest, and particularly on the heart. In particular, heart and coronary artery injuries due to penetrating injury are uncommon compared to other organ damage, such as lung injury. However, our approach can be generalized to other anatomic regions by adding the necessary knowledge to the supporting ontologies. The FMA already includes comprehensive anatomic knowledge in the trunk, and the Direct Injury Reasoner could be extended to identify penetrating injury to other parts of the trunk. In a similar manner, the Secondary Injury Reasoner could be extended to reason about other types of injuries by creating the appropriate ontology class definitions and assertions. For example, the pleura of the lung could be defined as a vital structure, and the uninjured pleural space could be defined as absence of air in this space. A puncture of the pleura could be defined as allowing air into the pleural space; then a penetrating injury of the chest wall could be recognized to produce pneumothorax using automatic classification with this ontology.

4. Conclusion

We have demonstrated methods of augmenting spatial geometric models of injured subjects with anatomic knowledge sources in ontologies to develop intelligent reasoning services. Detailed anatomic knowledge, usually only available to an expert, can be encoded in ontologies and exploited by computer applications to reason about the consequences of penetrating injury. The knowledge used by these applications is in a declarative format that can be maintained by domain experts and interpreted by machines. These knowledge-based tools may help reduce the complexity and potential confusion associated with assessing these injuries.

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. We are grateful to Cornelius Rosse for many years of exciting discussion and collaboration. This work was also supported by the Protégé resource, under grant LM007885 from the U.S. National Library of Medicine.

References

[1] West JG, Trunkey DD, Lim RC. Systems of trauma care. A study of two counties. Arch Surg 1979;114(4):455—60.
[2] Virtual Soldier Project. http://www.virtualsoldier.net/summary.htm, 2004. Accessed March 21, 2006.

[3] Rosse C, Mejino Jr JL. A reference ontology for biomedical informatics: the Foundational Model of Anatomy. J Biomed Inform 2003;36(6):478—500.
[4] Noy NF, Musen MA, Mejino JLV, Rosse C. Pushing the envelope: challenges in a frame-based representation of human anatomy. Data Knowl Eng 2004;48(3):335—59.
[5] Smith MK, Welty C, McGuinness D. OWL Web Ontology Language Guide, http://www.w3.org/TR/owl-guide/, 2004. Accessed March 21, 2006.
[6] Ackerman MJ, Yoo TS. The Visible Human Data Sets (VHD) and Insight Toolkit (ITk): Experiments in Open Source Software. Proc AMIA Symp 2003;773.
[7] Lu K, Clark Jr JW, Ghorbel FH, Ware DL, Bidani A. A human cardiopulmonary system model applied to the analysis of the Valsalva maneuver. Am J Physiol Heart Circ Physiol 2001; 281(6):H2661—79.
[8] Bassingthwaighte JB. Back to fundamentals: anatomy-based physiological bioengineering. Ann Biomed Eng 2000;28(7): 701—3.
[9] Freeman KM, Thompson SF, Allely EB, Sobel AL, Stansfield SA, Pugh WM. A virtual reality patient simulation system for teaching emergency response skills to U.S. Navy medical providers. Prehospital Disaster Med 2001; 16(1):3—8.
[10] Ogunyemi OI, Clarke JR, Ash N, Webber BL. Combining geometric and probabilistic reasoning for computer-based penetrating-trauma assessment. J Am Med Inform Assoc 2002;9(3):273—82.