

Supporting the Development of Medical Ontologies

David CORSAR^a, Laura MOSS^a, Derek SLEEMAN^{a,b} Malcolm SIM^b

^a *Department of Computing Science, University of Aberdeen, Aberdeen, UK*

^b *Section of Anaesthesia, Glasgow Royal Infirmary, University of Glasgow, Glasgow, UK*

Abstract. Ontologies are widely used in the biomedical community, which has built standard reference ontologies for various aspects of medicine. These projects have produced broad descriptions of the medical domain, resulting in large, complex ontologies which can be difficult to reuse as part of a focused application. We describe four ontologies which we developed to support intelligent reasoning about a particular medical sub-domain. We also describe how concepts in these ontologies can be aligned with standard reference ontologies to promote interoperability.

Keywords. Intensive Care Unit, ICU, Ontology, Reasoning

Introduction

Ontologies have recently become “*institutionalized*” in the biomedical community [23], serving a range of purposes from knowledge sources in specific, single purpose applications built by one research group, through to providing wide ranging descriptions of domains built by international collaborations (discussed in section 1). These projects have exploited the purpose of ontologies which according to Gruber’s definition is to create an “*explicit specification of a conceptualization*” [6], thereby providing the community with a common representation and vocabulary for describing and analysing data [23]. However, as also noted in [23], these ontologies have become large and complex, which hinders both their manageability from a maintenance viewpoint, and their (re-)use as components of systems.

One technique gaining attention from the Semantic Web community for supporting the maintenance, reuse, and evolution of large ontologies is modularisation. One popular approach to ontology modularisation is the extraction of significant modules of related knowledge (required for a particular use) from existing, large ontologies [13,18,4,2,5]; however, the variety and complexity of proposed approaches have been criticised for hampering the reuse process [3].

We recently needed to rapidly build an ontology/ontologies in conjunction with clinicians from Glasgow Royal Infirmary which was capable of supporting a system reasoning about Intensive Care Unit (ICU) time series data. The ontologies were required to provide knowledge about areas such as disorders, treatments, and patient data. Ideally we would have reused (parts of) existing relevant ontologies; however, although various ICU ontologies have been discussed in the literature (including [9,15,12]), none were

available to us at that time; so the primary option considered was to use a module extraction technique to produce modules containing appropriate concepts from the various standard biomedical ontologies. Several problems were identified with this approach:

- A range of module extraction techniques exist, which would have required significant analysis, and understanding. To our knowledge to date there is no guide as to which extraction method is most effective in particular situations.
- It can be difficult to identify the concepts relevant to our planned system in the appropriate biomedical ontologies which contain thousands of concepts.
- Module extraction techniques typically focus on extracting all the concepts related to particular concepts; this could result in modules consisting of concepts we did not require.
- Having extracted a series of modules, relationships would need to be defined between these modules as appropriate for our application.
- There is no guarantee that after performing these steps, the resulting ontologies would contain all the knowledge required by our system.
- This process would have to be repeated every time a new concept was required by the system.

In the information systems context, the standard biomedical ontologies would be referred to as “domain ontologies” [7], that is ontologies provide the vocabulary related to a generic domain (such as medicine). Our system requires the development of “application ontologies” [7], which specialise the concept descriptions in the domain ontologies and further “task ontologies” (which describe a particular type of task, such as diagnosis) to produce concept descriptions specifically for an application. As our development team was small (one person), and rapid development of the system was vital, we decided not to use the reuse-based process outlined above, instead we opted to build our own (relatively small) ontologies for the task. This produced compact ontologies, consisting solely of the knowledge required by the system. This approach obviously has many benefits when: building, using, and editing the ontologies with standard ontology editors (e.g. Protégé 3), which do not handle large ontologies well; maintaining the ontologies is easier, simply extending the relevant ontology when needed; and inferencing using these ontologies is quicker (important for the system’s end user). We do, however, recognise that standardised biomedical ontologies have a role as an important reference point, particularly beneficial for supporting interoperability, and so should be used whenever possible. To achieve this, we developed a simple alignment meta-ontology which enables the definition of correspondences between concepts in our ontologies with concepts in other ontologies. We believe this approach is consistent with the description of FMA and SNOMED-CT as reference terminologies/ontologies [20,21].

We have since analysed the ontologies we developed, generalising and adapting them to produce, we believe, three ontologies which provide a framework for describing the concepts required to perform various types of knowledge-based reasoning in a medical (sub-)domain. We believe that these generic descriptions can be quickly extended to describe a particular medical domain, and so can be used to support the rapid development of knowledge-based systems. As the extended models can be aligned with appropriate standard medical ontologies/knowledge sources, our approach supports easy integration and use of different standard biomedical ontologies by an application. We believe

that our ontologies can be viewed as an initial step towards providing an ontology design pattern for developing medical application ontologies.

This paper is structured as follows: section 1 summarises existing biomedical ontologies; section 2 outlines our ontologies for alignment, human physiology, and medical domains; section 3 describes the use of these ontologies when developing applications; section 4 outlines an existing application using our ontologies; and section 5 provides some conclusions and outlines future work plans.

1. Existing Biomedical Ontologies

This section provides an overview of four established biomedical ontologies¹, which provide definitions for various medical concepts for reuse or reference.

Unified Medical Language System (UMLS), which is produced and distributed by the US National Library of Medicine for the purpose of aiding “*the development of computer systems that behave as if they ‘understand’ the meaning of the language of biomedicine and health*”[1]. The UMLS Semantic Network consists of subject categories (derived from many different thesauri, classifications and controlled terms) and defines relationships between these categories. The major groups of categories are: organisms, anatomical structures, biological function, chemicals, events, physical objects, and ideas [1]. The UMLS Semantic Network is provided in a relational table format.

SNOMED Clinical Terms (SNOMED-CT) [21] is a widely-used clinical terminology that provides a common reference point for the comparison and aggregation of data. Concepts are organized into a hierarchy from the general to the specific and include the following high level categories: clinical finding/disorder, procedure/interventions, observable entities, body structure, organism, substance, pharmaceutical/biological product, specimen, special concept, physical object, physical force, event, environmental/geographic location, social context, and staging/scales. SNOMED-CT currently contains more than 311,000 concepts; each concept can be associated with a concept description; SNOMED-CT contains 800,000 concept descriptions and approximately 1,360,000 relationships between concepts.

The Foundational Model of Anatomy (FMA) Ontology is a comprehensive collection of classes and relationships which describe the structural organisation of the human body from the molecular level to major body parts. The FMA is implemented in a frame-based formalism and stored in a relational database; it contains 70,000 distinct anatomical concepts, 110,000 terms, and 170 relationships [17]. The proposed use of the FMA is as a reference ontology for “*correlating different views on anatomy, aligning existing and emerging ontologies in bioinformatics ontologies*”.

GALEN Common Reference Model The GALEN Common Reference Model (part of the Generalized Architecture for Languages, Encyclopaedias and Nomenclatures in medicine (GALEN) project) contains a large open-source ontology for the medical domain formulated in a specialised description logic based language called GRAIL [16]. The GALEN Common Reference Model represents the core concepts in pathology, anatomy and therapeutics, that have widespread use in medical applications.

¹It is acknowledged that there are numerous other examples of biomedical ontologies.




Figure 1. Visualisation of the Alignment meta-ontology.

2. A Framework for the Development of Medical Ontologies

As discussed in the introduction, using standard biomedical ontologies for specific applications is not easy, mainly due to their large size. This characteristic is to be expected - the goal of these projects is to supply knowledge for a (wide) range of applications, and so they require general models, which consequently become very large and complex. The task therefore becomes that of harnessing the power of these large general ontologies while still building applications quickly. To support this task, we outline three ontologies (expressed in OWL-DL²) which can be used as a base for models of medical domains used in knowledge-based systems; further these models can be aligned with other knowledge sources to support interoperability through reuse of general medical ontologies.

2.1. Alignment Annotations

The alignment meta-ontology (visualised in figure 1) defines a `hasAlignment` annotation object property, which can be used to define a correspondence between a concept in an ontology and a concept in another knowledge source. The range of the `hasAlignment` property is the `Alignment` class, which refers to the other knowledge source (via the `inSource` property) and the name of the corresponding concept in that other source (via the `conceptName` property). The `Knowledge_Source` class stores both the URI and the name of a knowledge source. This allows alignments to be made to other knowledge sources of any type, such as ontologies (for example, GALEN), semantic nets (for example, UMLS), thesauri (for example WordNet), or Web pages (for example, the NCI online browser for SNOMED-CT³). In cases where the knowledge source is not an ontology, the URI should be set to either the local file or remote/Web location.

For example, an alignment associated with a `Steroid` class may be with the “Steroid (substance)” concept in SNOMED-CT. The SNOMED-CT knowledge source would be represented as an individual of type `Knowledge_Source`, with `name` property value “SNOMED-CT” and an appropriate `uri` property value; the alignment would be represented by an individual of type `Alignment`, with the `conceptName` property value of “Steroid (substance)” and `inSource` property value of the SNOMED-CT knowledge source individual; this `Alignment` individual would then be set to the value of the `hasAlignment` annotation property for the `Steroid` class. The alignment ontology is imported by the three ontologies described below.

²<http://www.w3.org/TR/owl-features/>

³<http://nciterns.nci.nih.gov/NCIBrowser>

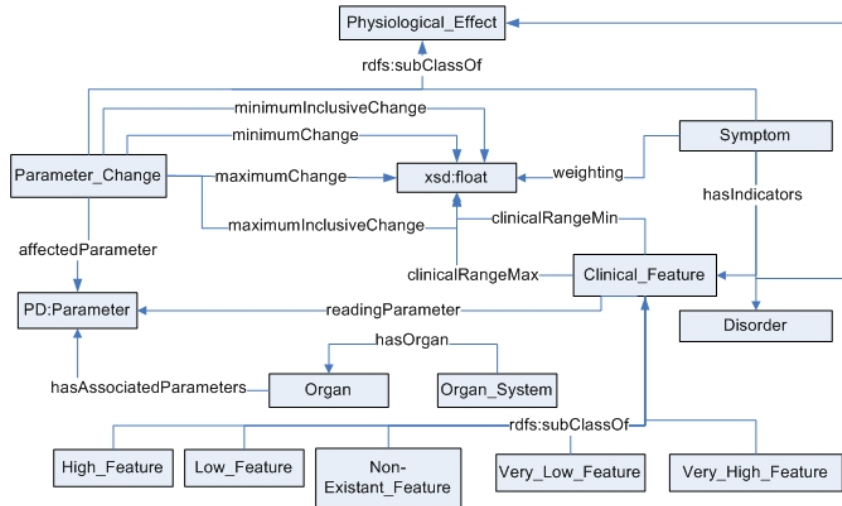


Figure 2. Visualisation of the Human Physiology Ontology.

2.2. The Human Physiology Ontology

Knowledge of human physiology is often required when reasoning in the medical domain, and so we have defined a basic human physiology ontology (visualised in figure 2) to represent knowledge of organs and organ systems, clinical features, and physiological effects.

2.2.1. Describing Organs and Systems

The `Organ_System` and `Organ` classes are used to represent basic human physiology, with the `hasOrgan` property associating the organ system with its primary organ. For example, the “respiratory system” (which could be represented as an individual of type `Organ_System`) has the “lungs” (which could be represented as an individual of type `Organ`) as the primary organ. The `hasAssociatedParameters` is used to associate `Physiological_Parameters` with an `Organ`. For example, the parameters “ FiO_2 ” and “ SpO_2 ” are physiological parameters associated with the “lungs”.

2.2.2. Clinical Features and Physiological Effects

The `Clinical_Feature` class provides a template for describing physiological states such as ‘*low MAP*’ and ‘*high temperature*’. The various subclasses of `Clinical_Feature` each describe a different level/type of feature, such as a very low, low, high, or very high value. For example ‘*low MAP*’ is an instance of `Low_Feature` associated with the parameter MAP with a value in the range 51 to 69.

The `Physiological_Effect` class represents different types of effects that occur in the human body; two types of effects have been defined: parameter changes and symptoms.

The `Parameter_Change` class describes changes in physiological parameters; a string description of the change is provided, and the magnitude of the change

(via the `minimumChange` and `maximumChange` properties). The three subclasses of `Parameter_Change` define the possible directions of the change.

The `Symptom` class represents the physiological symptoms associated with a disorder. A symptom can have various types of indicators: physiological effect(s), clinical feature(s), and/or disorder(s). As there may be more than one symptom associated with a disorder, a `weighting` can be applied to each symptom to reflect its overall importance to a disorder. Subclasses of `Symptoms` can be used to reflect the different types of symptom.

2.3. The Medical Domain Ontology

The medical domain ontology provides a system with the knowledge specific to a particular medical specialty. Four aspects of a particular medical domain can be described: disorders, treatments, disorder severity scores, and drugs.

2.3.1. Describing Disorders

The main purpose of the `Disorders` class is to provide information about commonly encountered medical conditions in a particular medical domain. Depending on the domain and the application in which the ontology is being used, it may be appropriate to describe various types of disorders; this can be achieved by defining subclasses of the `Disorder` class. Currently there is one such subclass, `Medical_Disorders`, which has eleven different direct subclasses representing disorders of different physiological systems, each with appropriate subclasses.

The properties `requiredSymptoms` and `additionalSymptoms` provide information about the clinical signs and symptoms associated with a particular disorder. All of the required symptoms must be observed before a diagnosis can be made; observation of any *additional* symptoms strengthens the diagnosis. Casual relationships between different medical disorders are frequently observed in some domains, and can be described using the `causedBy` and `causes` (inverse) properties. This allows one to define relationships such as the disorder ‘pulmonary oedema’ which can lead to the disorder ‘heart failure’. Finally, the treatments used for a particular disorder are associated with the disorder via the `treatments` property.

2.3.2. Describing Treatments

The `Treatment` class represents the commonly used treatments in a particular medical domain. There are many different types of treatments across the whole of medicine, and we do not attempt to model each one; rather we expect developers to define subclasses of `Treatment` as appropriate for the application they are developing. For example, when building our ICU ontology, we defined four types (subclasses) of `Treatment`: `Drug`, `Procedures`, `Nutrition`, and `Fluids`. In this implementation, the definition of the `Treatment` class is minimal, with only the `forDisorder` property associated with it in order to link a treatment with a disorder.

2.3.3. Describing Disorder Severity

Scoring systems are often used to summarise a patient’s condition (for example, abstract labels such as ‘A’, ‘B’, ‘C’, or standard scoring systems such as Apache IV [8] or Glas-

gow Coma Score [22]); scoring systems can include a definition of ‘*acceptable*’ physiological ranges. We represent scoring systems using the `Scoring_Stage` class in the medical domain ontology.

The `Scoring_Stage` class is associated with a `label` property, which has the range `Rating`, this is used to define a particular score/label (which could be, for example “acceptable”). The `Scoring_Stage` class is also associated with the `parameterRange` class, which has the range `Score_Element`. A `Score_Element` essentially defines a combination of parameter ranges, which (overall) must evaluate to true for the patient to be considered in that particular `Scoring_Stage`. For example, a rating ‘B’ in the Glasgow scoring system [19] may require (amongst other parameters) that the patient’s oxygen saturation (SpO_2) be between 94% and 95%.

2.3.4. Describing Drugs and Related Details

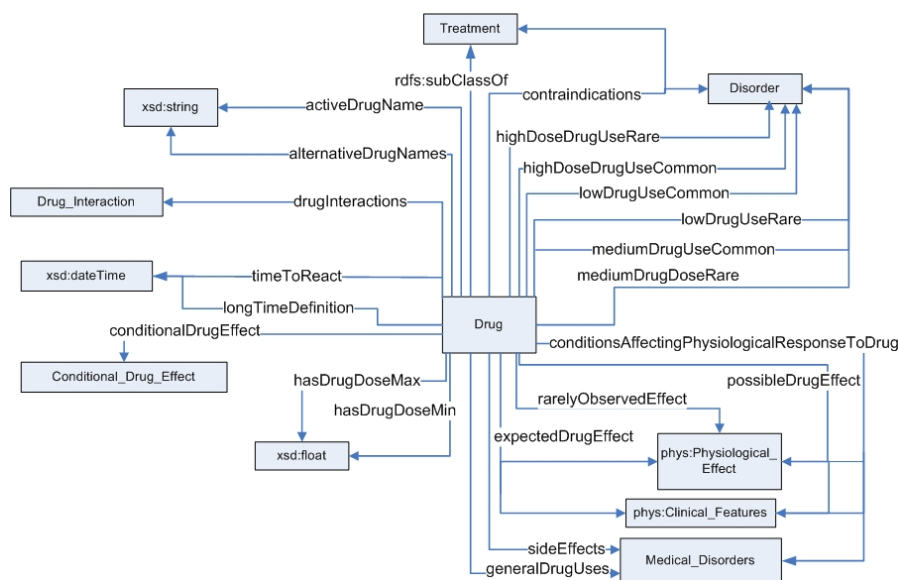


Figure 3. Visualisation of the drug class from the Medical Domain Ontology.

We have used the `Drug` class (a subclass of `Treatment`) (visualised in figure 3) to describe how drugs are used as treatments in the ICU domain, although we believe the class description is generic and can be applied when modelling other medical domains. Features of a drug, such as `activeDrugName`, `alternativeDrugNames`, the anticipated length of time between the drug’s administration to a patient and its effect being observed (the drug’s `timeToReact`) and any contraindications of the drug (i.e. known scenarios (disorders or other treatments) in which a patient should not be given a particular drug) have been modelled in the drug class; in addition, descriptions of a drug’s effects, interactions, and uses are supported.

Various types of drug effects are described: *expected* effects; *possible* effects, which are effects that may occur; *conditional* effects that occur under certain conditions, for example, when the drug adrenaline is administered it would *not* normally be expected

that a patient's SpO₂ would increase, however, when the patient also has a low cardiac output an increase in SpO₂ would be observed; *rarely observed* effects are very unusual but still theoretically possible; and (unwanted) *side* effects. In general, an effect is either a physiological effect or a clinical feature.

When drugs are given to a patient they can interact with other drugs to produce anticipated and unanticipated physiological effects, and so having knowledge of drug interactions can be important when reasoning in a medical domain. Due to the uncertain nature of drug interactions, they are, in general, difficult to model in OWL ontologies, however we have developed a specific definition of an interaction that would support the reasoning we required. For a particular drug, it is usually known which other drugs can interact with it and the associated nature of the interaction. This is represented by the Drug_Interaction class. Each Drug_Interaction is associated with a drug (hasDrug), the interacting drug (interactsWith), and the physiological effects observed during the interaction (interactionEffect).

Occasionally a physiological condition may affect the anticipated response to a drug, and this can be specified using the conditionsAffectingPhysiologicalResponseToDrug property. For example, severe sepsis can result in a patient being unable to respond to the drug norepinephrine. The condition which affects the drug can take the form of a clinical feature, e.g. high MAP; physiological effect, e.g. increase in FiO₂⁴; or another medical disorder, e.g. sepsis.

Various properties are used to represent different drug doses, as an individual drug may be given at different doses depending on the severity of the disorder. How commonly a drug is used to treat particular disorders can also be described.

2.4. The Patient Data ontology

The Patient Data ontology (visualised in figure 4) has been designed to model the time series data which is typically collected in medical domains. The ontology defines a Patient_Data class, which represents the patient's Sessions and Location. The Session class models a treatment session, which in turn links a series of Timepoints; the later describe the Readings for a particular xsd:dateTime. Each Reading has a Parameter and a value.

In our experience the vast majority of the information used by clinicians when treating ICU patients is the data collected by the unit. The exceptions being that lab test data

⁴Fraction of inspired oxygen given to the patient.




Figure 4. Visualisation of the Patient Data Ontology.

and sometimes a diagnosis (reason for patient being in the ICU) are included; such data could be easily recorded through extensions to the patient data ontology.

3. Supporting Application Development/Research

We believe the ontologies described above provide a framework for specifying the types of knowledge a knowledge-based system requires to perform reasoning within a medical domain. However, to enable a system to perform such reasoning in a particular medical sub-domain, these ontologies must be extended and/or instantiated to describe the appropriate domain knowledge. We envisage that this framework will enable the rapid development of a knowledge-based system within a medical domain in the following manner:

1. The knowledge engineer becomes familiar with the types of medical domain knowledge that the system requires.
2. The knowledge engineer uses this knowledge to extend our medical and physiological ontologies with further class definitions.
3. The extended ontologies are then instantiated, providing further domain knowledge for the system⁵.
4. Once the knowledge engineer is satisfied with their ontologies, the alignment annotation property can be used to define correspondences between the “local” concepts and those in other knowledge sources (for example, the standard reference ontologies discussed in section 1).
5. If necessary, the patient data ontology is instantiated with data from a medical source. This would normally be performed by conversion software which performs this task at runtime.
6. Having produced a suitable domain model (aligned with other knowledge sources) and patient data, a system can be produced using that knowledge.
7. If necessary, this process (particularly steps 1-4) can be repeated at any time to extend the domain knowledge as required.

4. Example Application - Explaining Anomalous Patient Responses to Treatment in the Intensive Care Unit

Within the medical domain there are clear expectations of how a patient should respond to treatments administered. When these responses are not observed, this can indicate a serious condition for the patient, and can also be challenging for clinicians to understand why the anomalous responses have occurred. An ontology-driven system implemented for the Intensive Care Unit (ICU) domain is currently being developed to assist clinicians detect anomalous patient responses to treatment and suggesting hypotheses to explain these anomalies. As part of the development of this tool, we have undertaken a study to determine how ICU clinicians identify anomalous patient responses; we then asked further clinicians to provide potential explanations for these anomalies [10], and analysed the high-level reasoning deployed by the clinicians when determining potential explana-

⁵We appreciate that it is likely the first three steps will be performed recursively as the knowledge engineer's understanding of the domain (and consequently, the ontologies) evolve.

tions. Figure 5 provides an example of an anomaly identified by a clinician and figure 6 provides an example of a hypothesis provided in response to the anomaly shown in figure 5.

“...but then we obviously do something because the cardiac output and the cardiac index get a bit better and the thing that we seem to have done is put the noradrenaline up to a high dose, but that isn't necessarily quite what we would expect from a high dose of noradrenaline”

Figure 5. An example of an identified anomalous response to treatment.

“The only thing that I can think of is that noradrenaline is actually an inotrope. In a low dose, it tends to be a vasoconstrictor, in higher doses it's an inotrope. So it might just be that, that dose for that particular patient is enough to, as well as causing a tightening, is enough to cause an increase force of contraction as well”

Figure 6. An example of an explanation for an anomaly.

To summarise, during the interviews the clinicians appeared to use domain knowledge about treatments, medical conditions, and the desired physiological state of the patient to explain a given anomalous treatment or effect. Domain knowledge was also applied by the clinicians whilst examining the patient's data to determine new facts, for example, that the patient is suffering from a myocardial infarction, and also to test/eliminate potential hypotheses.

This high level reasoning demonstrated by the clinicians was generalised and used as the basis for strategies implemented in a system to suggest hypotheses for anomalous patient responses to treatment. This part of the system uses a knowledge base consisting of the OWL ontologies described previously which have been extended and instantiated for the ICU domain and a Java based program implementing the hypothesis generation strategies extracted from the domain experts' protocols.

Currently, 11 different hypothesis generation strategies have been identified and implemented, which make use of the ICU ontologies and patient data in various ways. For example, the “drug” strategy attempts to explain an anomalous effect by identifying whether another drug that the patient is receiving at the time the anomaly was observed could have the same effect (side effect or expected effect) as the observed anomalous effect (and hence this further drug explains the observed anomaly). To achieve this, the strategy queries the ICU medical ontology for all the individuals of type `Drug` that have `sideEffects` or `expectedDrugEffect` equal to the observed anomalous effect. The patient data that corresponds to the time that the anomaly was observed is then queried to determine if the patient was taking any of the drugs (returned by the first query) at that time. If so, a hypothesis is returned to the user for each of the matching drugs suggesting that each drug could be responsible for the anomalous effect.

An initial evaluation of this hypothesis generation system has been performed using 15 test cases, each containing an individual anomaly. In total, the system produce 13 different hypotheses after running the test cases (an average of 0.9 hypotheses per test case), which were presented to and accepted by a senior ICU clinician.

The framework proposed in this paper has been successfully used as part of the hypothesis generation tool discussed in this section. It is intended that further applications and subsequent evaluations of the framework will be carried out during the development of a knowledge-based system in a different medical domain, namely the renal dialysis domain.

5. Conclusions and Future Work

As previously mentioned, ontologies are widely used in the biomedical domain; several standard reference ontologies for various aspects of biomedicine already exist. These ontologies, although very general, are also very large, which can inhibit their reuse in applications. We have outlined three small ontologies which provide a framework for describing medical concepts and medical domain models. We have also outlined a fourth, alignment ontology (section 2.1) which can be used to align concepts in the developed domain models with those in other knowledge sources; further, we discussed how these ontologies can be used in the development of a new system. Below we outline two aspects of future work using these ontologies/approach.

The Architecture for Clinical Hypothesis Evaluation (ACHE) [11] aims to support the performance of different studies using data from multiple sources. This is achieved by using a generic repository for time series medical data from a variety of sites, and a tool which supports transfer of data from a medical source to the repository. As part of the transfer process, users are required to define mappings between the parameter names in the incoming data and those in the ACHE repository; this could easily be updated to allow mappings to concepts stored in a medical domain ontology. A web service is provided for accessing and updating the repository; this could also be updated to use the patient data ontology for exchanging data, enabling any system using data expressed against that ontology to use data provided by ACHE. This would provide future systems with an easy method to acquire data in a supported format, regardless of the data's original format.

Knowledge-based systems often comprise of domain knowledge expressed as an ontology and the procedural (decision making) component expressed as a set of rules; in which case it is necessary to provide some linkage between the two components. The Semantic Web community have been working on formalisms for expressing rules against OWL ontologies, for example SWRL⁶. Currently, tools, such as the SwrJESSTab [14] plug-in for Protégé enable the user to define and run rules over instantiated ontologies, however they can be understandably slow when dealing with very large ontologies. Although extensions to our ontologies will likely result in sizeable ontologies, they should be considerably smaller than the available standard reference medical ontologies, and so current rule-based tools should run rules in a more acceptable time. We believe that using our ontologies with existing rule plug-ins will enable developers to quickly build and test rule-based systems in a variety of medical domains.

Acknowledgements

- Prof. John Kinsella (University of Glasgow) for helpful discussions.
- Kathryn Henderson and Jennifer McCallum (CareVue Project, Glasgow Royal Infirmary) & the staff and patients of the ICU Unit, Glasgow Royal Infirmary.
- This work was an extension of the routine audit process in Glasgow Royal Infirmary's ICU; requirements for further Ethical Committee Approval has been waved.
- This work has been supported under the EPSRC's grant number GR/N15764.

References

- [1] Unified Medical Language System (UMLS). <http://www.nlm.nih.gov/pubs/factsheets/umls.html>.

⁶<http://www.w3.org/Submission/SWRL/>

- [2] B. Cuenca-Grau, I. Horrocks, Y. Kazakov, and U. Sattler. Just the right amount: Extracting modules from ontologies. In *Proceedings of the 16th International World Wide Conference (WWW 2007)*, 2007.
- [3] M. d'Aquin, P. Doran, E. Motta, and V. Tamma. Towards a Parametric Ontology Modularization Framework Based on Graph Theory. In *Proceedings of the Second International Workshop on Modular Ontologies (WOMO 2007)*, 2007.
- [4] M. d'Aquin, M. Sabou, and E. Motta. Modularization: a Key for the Dynamic Selection of Relevant Knowledge Components. In *Proceedings of the ISWC 2006 Workshop on Modular Ontologies*, 2006.
- [5] P. Doran, V. Tamma, and L. Iannone. Ontology module extraction for ontology reuse: An ontology engineering perspective. In *Proceedings of the 2007 ACM CIKM International Conference on Information and Knowledge Management*, 2007.
- [6] T. Gruber. A Translation Approach to Portable Ontology Specification. *Knowledge Acquisition*, 5(2):199–220, 1993.
- [7] N. Guarino. Formal Ontology in Information Systems. In N. Guarino, editor, *Proceedings of FOIS'98*, pages 3–15. IOS Press, June 1998.
- [8] J. E. Zimmerman MD FCCM, A. A. Kramer PhD, D. S. McNair MD PhD, F. M. Malila RN MS. Acute Physiology and Chronic Health Evaluation (APACHE) IV: Hospital mortality assessment for today's critically ill patients. *Critical Care Medicine*, 34:5:1297–1311, 2006.
- [9] J. M. Juarez, M. Campos, J. Palma, and R. Marin. Computing context-dependent temporal diagnosis in complex domains. *Expert Systems with Applications*, 35:991–1010, 2008.
- [10] L. Moss, D. Sleeman, M. Booth, M. Daniel, L. Donaldson, C. Gilhooly, M. Hughes, M. Sim, and J. Kinsella. Explaining Anomalous Responses to Treatment in the Intensive Care Unit. In *Proceedings of 12th Conference on Artificial Intelligence in Medicine (AIME'09)*, 2009.
- [11] L. Moss, D. Sleeman, J. Kinsella, and M. Sim. ACHE: an Architecture for Clinical Hypothesis Examination. In *Proceedings of 21st IEEE Symposium on Computer-Based Medical Systems (CBMS 2008)*, 2008.
- [12] N.F de Keizer, A. Abu-Hanna, R. Cornet, J.H.M. Zwetsloot-Schonk, C.P. Stoutenbeek. Analysis and Design of an Ontology for Intensive Care Diagnoses. *Methods of Information in Medicine*, 38:102–112, 1999.
- [13] N. Noy and M. Musen. Specifying Ontology Views by Traversal. In *Proceedings of the International Semantic Web Conference (ISWC)*, 2004.
- [14] M. J. O'Connor, H. Knublauch, S. W. Tu, B. Grosso, M. Dean, W. E. Grosso, and M. A. Musen. Supporting Rule System Interoperability on the Semantic Web with SWRL. In *Proceedings of the Fourth International Semantic Web Conference (ISWC2005)*, pages 974–986. Springer, November 6–10 2005.
- [15] J. Patrick, Y. Wang, P. Budd, A. Rector, S. Brandt, J. Rogers, R. Herkes, A. Ryan, and B. Vazirnezhad. Developing SNOMED CT Subsets from Clinical Notes for Intensive Care Service. In *Proceedings of Health Informatics New Zealand Conference (HINZ) 2008*, 2008.
- [16] A. L. Rector and W.A Nowlan. The GALEN project. *Computer Methods and Programs in Biomedicine*, 45:75–8, 1994.
- [17] C. Rosse and J. L.V. Mejino Jr. A reference ontology for biomedical informatics: the Foundational Model of Anatomy. *Journal of Biomedical Informatics*, 36:478–500, 2003.
- [18] J. Seidenberg and A. Rector. Web Ontology Segmentation: Analysis, Classification and Use. In *Proceedings of the World Wide Web Conference (WWW)*, 2006.
- [19] D. Sleeman, A. Aiken, L. Moss, J. Kinsella, and M. Sim. A system to detect inconsistencies between a domain expert's different perspectives on (classification) tasks. *Studies in Computational Intelligence: Recent Advances in Machine Learning*, In press.
- [20] B. Smith and C. Rosse. The role of foundational relations in the alignment of biomedical ontologies. *Medinfo*, 107:444–8, 2004.
- [21] K. A. Spackman and K. E. Campbell. SNOMED RT: A Reference Terminology for Health Care. *Proc. of the AMIA Annual Fall Symposium*, pages 640–644, 2007.
- [22] Teasdale, G and Jennett B. Assessment of coma and impaired consciousness. A practical scale. *Lancet*, 2:81–84, 1974.
- [23] A.C. Yu. Methods in biomedical ontology. *Journal of Biomedical Informatics*, 39(3):252–266, 2006.