

Hyperontology for the Biomedical Ontologist: A Sketch and Some Examples

Oliver Kutz¹, Till Mossakowski^{1,2}, Janna Hastings^{3,4}, Alexander Garcia Castro⁵, Aleksandra Sojic⁶

¹Research Center on Spatial Cognition, University of Bremen, Germany

²DFKI GmbH Bremen and University of Bremen, Germany

³Chemoinformatics and Metabolism, European Bioinformatics Institute, Cambridge, UK

⁴Swiss Center for Affective Sciences, University of Geneva, Switzerland

⁵University of Arkansas for Medical Sciences, USA

⁶European School of Molecular Medicine, Milan and University of Milan, Italy

Abstract. The Hyperontology framework has been recently introduced to provide a general methodology for heterogeneous ontology design, i.e. the construction of ontologies that have parts, or modules, written in different formalisms, and which are interlinked in complex ways. We here present a brief outline of this framework, discuss its features and merits, and illustrate its usefulness for the domain of biomedical ontology design by providing and discussing a number of examples.

1 Introduction and Motivation

Ontologies are today being applied in virtually all information-rich application areas, and in particular are of increasing importance in the Life Sciences [16].

While the OWL standard [18] has led to an important unification of notation and semantics, still many diverse formalisms are used for writing ontologies. Some of these, such as RDF, OBO [20] and UML, can be seen more or less as fragments and notational variants of OWL, while others, like F-logic and Common Logic, clearly go beyond the expressiveness of OWL. Moreover, not only the underlying logics are different, but also the modularity and structuring constructs, and the reasoning methods.

Many (domain) ontologies are written in description logics (DLs) such as *SROIQ(D)* (underlying OWL 2 DL) and its fragments. These logics are characterised by having a rather fine-tuned expressivity, exhibiting (still) decidable satisfiability problems, whilst being amenable to highly optimised implementations. However, there are many cases where either weaker DLs are enough – such as sub-Boolean *EL* (an OWL ‘profile’) – and more specialised (and faster) algorithms can be employed, or, contrarily, the expressivity has to be extended beyond the scope of standard DLs.

For example, a weaker DL suffices for the NCI thesaurus (containing about 45,000

concepts) which is intended to become the reference terminology for cancer research [17], but beyond DL expressivity is required for many foundational ontologies, for instance DOLCE [6], BFO¹, or GFO². Note however that these foundational ontologies also come in different versions ranging in expressivity, typically between OWL and first-order or even second-order logic (see Section 3.3 for a discussion of the kinds of problems this entails).

While the web ontology language OWL has evolved and extensions are being constantly developed, its main target application is the Semantic Web and related areas, and it can thus not be expected to be fit for any purpose: there will always be new, typically interdisciplinary application areas for ontologies where the employed (or required) formal languages do not directly fit into the OWL landscape. Heterogeneity (of ontology languages) is thus clearly an important issue. This does not only include cases where the expressivity of OWL is simply exceeded (such as when moving to full first-order logic), but, ultimately, also cases where combinations with or connections to formalism with different semantics have to be covered, such as temporal, spatial, or epistemic logics.

Biomedical ontologies in particular face the problem of heterogeneity, as the information

¹ See <http://www.ifomis.org/bfo/>

² See <http://www.onto-med.de/ontologies/gfo>

that is relevant for such ontologies comprises different data sources such as clinical and experimental data from various epistemic settings. For example, we consider the domain of diseases. A patient's information might include age, family history of disease and social status on the one hand, and on the other hand experimental data for that patient might include metabolic profiles, tumour and genetic markers. Therefore, ontologies of disease need to stretch from an epidemiological, through a traditional clinical representation, to the ontology that includes specific molecular pathways and interactions. In particular, ontologies for complex diseases such as cancer have to deal with spatio-temporal heterogeneity, combinations of qualitative and quantitative data, and missing links between physiological and pathological data [4].

Moreover, in biomedical domains many unknowns still remain, and the questions and theories that drive experimental research also shape the spatial and temporal boundaries of representation. For example, whether mitochondria are classified as organisms or as cellular components depends on the background theory that is accepted, i.e. whether mitochondria are prokaryotes living within eukaryotic cells. Thus, heterogeneity in the ontologies might originate not only from the different formalisms used, but also from heterogeneity within and across specific domains. Therefore, an ontology integration that is intolerant to ontological heterogeneity might not only be unfeasible in practice but also impossible in principle.

We here suggest a heterogeneous framework for the design of biomedical ontologies, based on the theory of hyperontologies as introduced in [12], which suggests solutions to some of these problems of heterogeneity. In particular:

- (i) We briefly sketch the main features of hyperontologies in Sec. 2, including reasoning support based on the tool HETS;
- (ii) We discuss how these features can in general be used within the context of biomedical ontologies in Sec. 3, focusing on 3 aspects in particular, namely (1) borrowing of tools, semantics, and reasoners via logic translations (2) structuring and modularity, and (3) (heterogeneous) ontology refinement.

- (iii) We finally present several examples from the world of biomedical ontology engineering in Sec. 4, illustrating how the structuring techniques for heterogeneous biomedical ontologies can be used in practice; Sec. 5 summarises and discusses future work.

2 A Very Brief Sketch of the Hyperontology Framework

In the presence of several alternative choices of modelling formalisms, it can be a rather difficult task for an ontology designer to choose an appropriate logic and formalism for a specific ontology design beforehand – and failing in making the right choice might lead to the necessity of re-designing large parts of an ontology from scratch, or limit future expandability. Another issue is the mere size of ontologies making the design process potentially quite hard and error prone (at least for humans), which is particularly a problem for ontologies in the Life Sciences. This issue has been partly cured in OWL by the `imports` construct, but still leaves the problem of ‘debugging’ large ontologies as an important issue, see e.g. [10]. Also, simple operations such as the re-use of parts of an ontology in a different ‘context’ whilst *renaming* (parts of) the signature are not possible in the OWL languages, making it difficult to combine ontologies that use the same terms analysed from different modelling perspectives, thereby easily yielding inconsistencies when performing naive ontology combination.

We here propose a solution to the above issues based on the concept of *heterogeneity*: facing the fact that several logics and formalisms are used for designing ontologies, we suggest heterogeneous structuring constructs that allow to combine ontologies in various ways, in a systematic and formally and semantically well-founded fashion. Our approach is based on the theory of institutions (which is a sort of abstract model theory) and formal structuring techniques from algebraic specification theory. Its main features are the following, paraphrasing [12]:³

- The ontology designer can use OBO or OWL

³ For technical detail and extensive discussion we have to refer to [12].

to specify most parts of an ontology, and can use first-order (or even higher-order) logic where needed. Moreover, the overall ontology can be assembled from (and can be split up into) semantically meaningful parts ('modules') that are systematically related by structuring mechanisms. These parts can then be re-used and/or extended in different settings.

- Institution theory provides the framework for formalizing 'logic translations' between different ontology languages, translating the syntax and semantics of different formalisms. These translations allow in particular the 'borrowing' of reasoning and editing tools from one logic to another, when appropriately related.
- Various concepts of 'ontological module' are covered, including simple imports (extensions) and union of theories, as well as conservative and definitional extensions.
- Structuring into modules is made explicit in the ontology and generates so-called proof obligations for conservativity. Proof obligations can also be used to keep track of desired consequences of an ontology, especially during the design process.
- Re-using (parts of) ontologies whilst renaming (parts of) the signature is handled by *symbol maps* and *hiding symbols*: essentially, this allows the internalisation of (strict) alignment mappings.
- The approach allows heterogeneous refinements: it is possible to prove that an ontology O_2 is a refinement of another ontology O_1 , formalised in a different logic. For instance, one can check if a domain ontology is a refinement of (a part of) a foundational one. An interesting by-product of the definition of heterogeneous refinements is that it also provides a rather general definition of heterogeneous sub-ontology and of ontology equivalence.

Tool support for developing heterogeneous ontologies is available via the Heterogeneous

Tool Set HETS, which provides parsing, static analysis and proof management for heterogeneous logical theories. HETS visualises the module structure of complex logical theories, using so-called development graphs. For individual nodes (corresponding to logical theories) in such a graph, the concept hierarchy can be displayed. Moreover, HETS is able to prove intended consequences of theories, prove refinements between theories, or demonstrate their consistency. This is done by integrating several first-order provers and model-finders (SPASS, DARWIN), the higher-order prover (ISABELLE), as well as DL reasoners like PELLET and FACT++.

3 What Hyperontology Can Do for Biomedical Ontologies

3.1 Borrowing Reasoning and Editing Tools via Logic Translation

[14] lays the foundation for a distributed ontology language (DOL), which will allow users to use their own preferred ontology formalism whilst becoming interoperable with other formalisms. At the heart of this approach is a graph of ontology languages and translations. In connection with HETS, this graph enables users to:

- relate ontologies that are written in different formalisms (e.g. prove that the OWL version of DOLCE is logically entailed by the first-order version);
- re-use ontology modules even if they have been formulated in different formalisms;
- re-use ontology tools like theorem provers and module extractors along translations

A detailed discussion of the various translational relationships between (almost) all known ontology languages can be found in [14]. We here concentrate on the languages of specific interest for biomedical ontologies, namely OBO, OWL and its profiles, first- and second-order logic, and F-Logic and Common Logic.

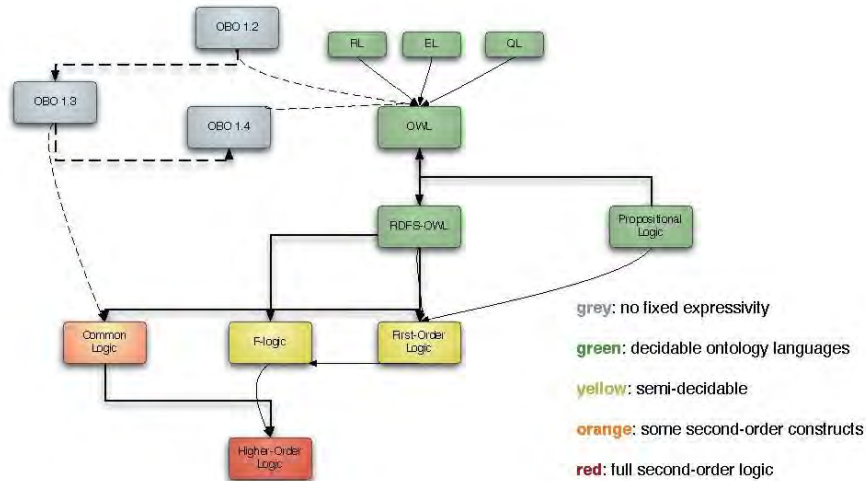


Figure 1. Translations between ontology languages

Fig. 1 illustrates the translational relationships. A ‘regular’ translation between two ontology languages, as marked by a solid arrow, means that the syntax and semantics of one logic can be translated into another. This means that, typically, the former is a fragment of the latter. A standard example would be OWL which, via the standard translation (which is available in HETS), can be considered a fragment of first-order logic. Note that translations concerning different versions of OBO are of different flavours⁴, thus are here marked by dashed lines. The OBO language does not itself come with formal semantics. Beginning with [7], who mapped a fragment of OBO 1.2 to OWL, a semantics for OBO has been assigned by translation. Version 1.3 of OBO, now abandoned, had something similar using Common Logic. The current specification of OBO, version 1.4, gets its semantics entirely via translation to OWL 2. In a sense, thus, the OBO language does not have a fixed logical expressivity, but depends on borrowed model-theoretic semantics from a particular mapping to another ontology language, relative to which corresponding reasoning methods and editing tools can be applied.

Logic translations can in particular be internalised in the ontology languages themselves, in the sense that ontologies can be written in a mix of logical formalisms, where

⁴ In particular, the progression between the different versions of OBO are only partial, leaving out some language constructs and adding others.

the translations assign respective semantics by operating in the background. For this to work properly, formal structuring principles are necessary, which we discuss next.

3.2 Structuring and Modularity

The web ontology language OWL as well as OBO can be accommodated within the larger framework of the heterogeneous common algebraic specification language HETCASL. Through this change in perspective, OWL and OBO can benefit from various useful HETCASL features concerning structuring, modularity, and heterogeneity. This tackles a major problem area in ontology engineering: re-use of ontologies and re-combination of ontological modules. We have briefly sketched the main structuring mechanisms already in the last section, namely unions and extensions of ontologies, translations along symbol maps, refinements, conservative extensions, etc.

To be able to write down such heterogeneous ontologies in a concise manner, we propose a structuring language that operates on top of and independently from a chosen ontology language. For instance, we use the notation `logic <logic-name>` to define the logic of the following specifications, which remains intact until that keyword occurs again. Similarly, an ontology (module) can be translated along a logic translation, which is written `<ontology>` with `logic <translation-name>`. The full language, which is also supported by HETS, cannot be given here, but compare [12].

3.3 Refinements: Relating Domain and Foundational Ontologies

Informally speaking, a (homogeneous) **refinement** of ontology O_1 into another ontology O_2 , both written in the same language, consists of a translation π which translates all of O_1 's axioms in such a way that the translated sentences follow from O_2 . For instance, a Biomedical domain ontology O written in OWL *refines* the OWL version of BFO exactly if O logically implies the translation of BFO's axioms.

But the approach also allows heterogeneous refinements: for instance, it is possible to prove that a first-order version BFO_{FOL} of a foundational ontology, here BFO, is a refinement of an OWL-based version BFO_{OWL} of BFO. Here, it needs to be established, using a first-order theorem prover, that all the translations of BFO_{OWL} 's axioms along the standard translation are logically entailed by BFO_{FOL} . Also, one can check if a domain ontology, written in OWL or OBO, is a refinement of (a part of) a foundational one, written in first-logic. This can be done by first *hiding* a part of the foundational signature, and then establishing a refinement. Note that hiding restricts the vocabulary of an ontology to an "export interface" (which is just a sub-vocabulary), while otherwise keeping the logical properties intact. All these verifications are supported by HETS.

4 Problems and Examples from Biomedical Ontology Design

4.1 Biomedical Imaging

When assessing the mechanical properties of bones, researchers use computational simulations to evaluate stress and strain maps under several boundary and load conditions. Such evaluations involve clinical data, e.g. pathological conditions of the patients, and mechanical properties of the materials to be used. Better models require high quality images, acquisition of which is not an easy task.

There are three main steps when doing computational simulations within the computational biomechanics domain. Pre-processing involves getting the geometrical model of the tissue; medical images obtained by methods

such as Scanning Electron Microscopy (SEM) or Microtomography (μ CT), are the main input for building these models. The images are thresholded, MIMICS can be used for this purpose; the quality of the model is directly related to the level of resolution and number of segmented images. The obtained CAD model offers sagittal, frontal and transversal planes; standard CAD software such as Inventor, Solid Edge, CATIA and Unigraphics are then used to manipulate the model. Finite Elements Methods (FEM) and the post-processing immediately follow; as our aim is to support sharing and reusability of medical images we are only focusing on the pre-processing phase. Details for pre-processing are illustrated in Fig. 2(a).

For describing a medical image it is often necessary to use several ontologies. For instance, Fig. 2(b) (left) illustrates the model for knee joints; hard tissue, e.g. Femur and Tibia bones, and soft tissue, e.g. Tibia and Femur cartilage, need to be identified. The osteoarthritis of the patient requires shaving the cartilage injury as presented in Fig. 2(b) (right). Such self-descriptiveness makes it possible for users to express complex queries such as: *'Retrieve knee joints images with cartilage injury.'*

Facilitating the execution of such a query requires the orchestration of several ontologies, namely: Radiology, Anatomy, Pathology, CAD. Ontologising the description of the DICOM file brings together radiological, spatial, anatomical and pathology related information; since these ontologies are space related, they are not necessarily available as OWL files, and therefore need to be heterogeneously combined using appropriate structuring mechanisms. The degree of segmentation together with other features will determine the suitability of the images. The report of the radiologist and the information contained in the DICOM should support the intelligent retrieval of information as well as the identification and management of anatomical features in the resulting CAD model. To support this with a tool chain is a challenge. Initial experiments in extracting both OWL ABox information and parameters for higher-order specifications from SolidWorks and CATIA CAD models and processing the result with HETS have been made in [11, 5].

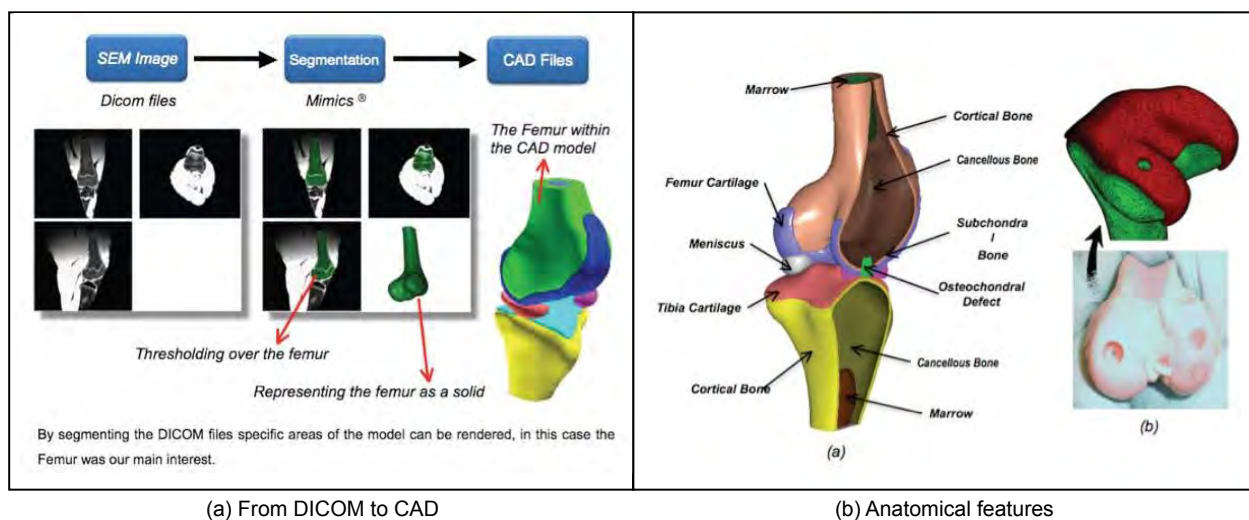


Figure 2. DICOM and Anatomy

4.2 Biochemical Structures

ChEBI (Chemical Entities of Biological Interest) is an ontology of chemical entities consisting of around 25000 entities in the latest release (April 2011) [3]. The core content of the ontology are molecules and ions which are biologically active in some fashion, whether naturally or artificially. The information encoded in the ontology includes a deep structural feature-based hierarchical classification for the chemical entities and a function-based encoding of the actions of the chemicals in biological contexts. For example, *morphine* (CHEBI:17303), an opiate analgesic drug, is included in the ChEBI ontology. It has structural classification, inter alia, in the classes *isoquinolines* and *alkaloids*, and function-based classification in classes *opioid analgesic* and *opioid receptor antagonist*.

Increasingly, OWL semantics is being used for the definition of classes of chemicals based on their shared structural features. Chemical structures are modelled as graphs in which atoms are the vertices and covalent bonds form the edges. For examples of this sort of approach, see [21] and [1]. This allows parts of chemical structures (such as, for example, a *carboxyl group*) to be used to fully define classes (such as, *carboxylic acids*). This is illustrated in Figure 3.

There are several challenges with this sort of approach. The first, well-known, is that chemical entities contain structural cycles. OWL is not suitable for modelling non-tree-like structures, and as a result other formalisms

must be used. Recent work has investigated the use of description graphs and rules to encode these structural features [8], but tool support for description graph-extended OWL ontologies is still poor. In the Hyperontology framework, we would be able to fully describe the structural features in a suitably expressive formalism, and link this to the core OWL ontology, without limiting which tools can be used. More concretely, parts of the modelling can be done in OWL, parts in first-order, and other parts e.g. in a suitable spatial logic, and reasoning is automatically delegated to an appropriate reasoner, possibly employing a logic translation to translate dependent OWL axioms to the first-order formalism.




Figure 3. Chemical structures used to define chemical classes

Facing similar challenges, the RNA ontology (RNAO) is an ontology for the structure and function of RNA molecules [9]. RNA molecules consist of chains of nucleotides which can display certain structural motifs or common

patterns. Encoding these structural motifs in the general case requires references to cyclic structures, which can be dealt with in rules or description graphs as for chemicals. Furthermore, the RNAO is provided in a first-order logical formalism implemented in SPASS, and a logically trimmed-down version implemented in OWL. The authors point out that giving a definition for an entirely covalently connected entity such as an RNA molecule based on atoms and bonds would require second-order logic in order to be properly formalised, and for this reason such a formalisation is not provided but the relevant features (transitive closure of the covalent connection relationship) are only approximated in the provided formalism. At present the different versions of the RNAO are not formally interlinked, and each has to be separately maintained and reasoned over. The hyperontology framework allows an elegant solution to this problem, allowing to formally relate different versions of the RNAO, e.g. by heterogeneous refinement (assuming the different versions are logically compatible), and to add second-order constraints on top of weaker formalisations.

4.3 Combining Bio-Ontologies

Since it provides a definition for all biologically interesting chemical entities, ChEBI aims to be sufficient for reuse, allowing use of its axioms as Lego bricks in defining specific molecules and molecular-entity-related biological entities. A quick inspection of ChEBI and a comparison to the related Lipid Ontology (henceforth LO) [1] reveals conceptual relations between the two ontologies. A more detailed review allows us to see the specialization of the axioms in the LO. However, these axioms are not always orthologous to those available in ChEBI, since the LO provides a far more detailed classification for lipids than is the concern of ChEBI. The hyperontology framework allows domain specifications such as the LO to effortlessly re-use parts of core ontologies such as ChEBI and even rename or redefine certain of their entities where needed. Also, more complex relationships between the ontologies' terms can be formalised in a heterogeneous ontology in the style of *Bridge Rules* as they are known from distributed DL or \mathcal{E} -Connections (see [12]).

ChEBI is also used as a reference in biological ontologies. Efforts are underway to explicitly link the Gene Ontology (GO) [19] to ChEBI through the OBO cross-product formalism [15]. Cross-products resemble OWL logical definitions composed in terms of intersection, that is, an 'and' operation. In the first example above, *1,3-dichloro-2-propanol metabolic process* would be formalised in the cross-product style as *metabolic process and has_participant some 1,3-dichloro-2-propanol*. Several different challenges arise in this ontology alignment process between ChEBI and GO. The first can be characterised as the problem of size explosion: ChEBI and GO are both upwards of 20000 terms with many more relationships, and as a result, reasoning over the combined ontologies can be prohibitively slow. The existing OWL:import mechanism requires the full content of both ontologies be loaded into an application (such as Protégé) in order to work with the cross-ontology definitions. The hyperontology framework allows us to bypass this problem with its built-in support for modularisation, even across ontology languages.

A further difficulty arises because of the common practice of classifying chemicals based on parts of the structure. Under this scheme, ChEBI's nucleotide is classified as a carbohydrate. However, in GO, a biological process such as nucleotide biosynthetic process is not considered to be a subtype of carbohydrate biosynthetic process. The straightforward combination of these two ontologies, ChEBI with its chemical-specific perspective, and GO with its biology-specific perspective, therefore leads to challenges for automated inferences. Present approaches to resolve such difficulties involve lengthy ongoing negotiations between the GO and the ChEBI editors to arrive at mutually satisfactory models that can be shared by both communities. A technological solution which allowed coexistence of the two perspectives without incorrect inferences in either would be better. The hyperontology framework's structuring and linking capabilities offer several roads towards this goal, of course making use of established methods such as statistics-based ontology matching.

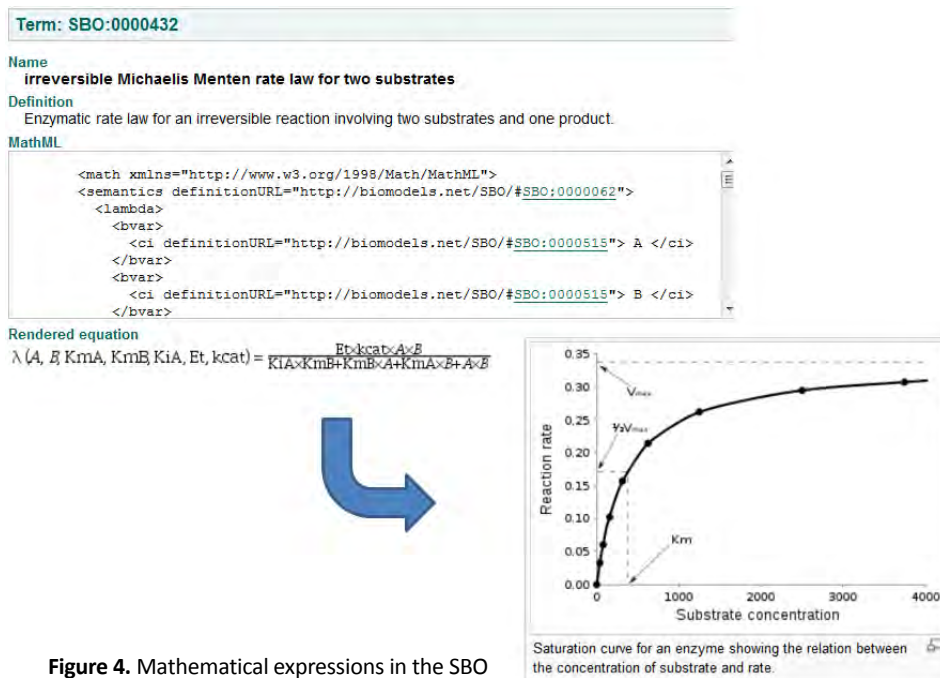


Figure 4. Mathematical expressions in the SBO

Another domain in which the integration of multiple ontologies is mandatory to the creation of successfully interoperable and reusable information is that of Systems Biology. Here, complex mathematical models are used to describe the behaviour of biological systems and to make predictions about their behaviour under different conditions. In order to exchange and unambiguously interpret such models, they need to be annotated with ontologies such as the Systems Biology Ontology (SBO) [13]. SBO contains many different types of entities: material entities such as proteins and small molecules; process participation roles such as inhibitor and stimulant; mathematical laws such as rate laws for biochemical reactions; and types of mathematical model experiments such as discrete and continuous, and many more besides. The hyperontology framework would allow a reformulation of the SBO as composed of modular units sourced from separate domain ontologies, a highly desirable goal. Furthermore, of particular interest in the SBO is that it captures complex mathematical relationships that can exist between biological entities in dynamic conditions. In SBO, these relationships are currently expressed in MathML. It is an open challenge to expose some of the relational information encoded in the SBO mathematical equations to ontological reasoning. This would require interrelating different formalisms,

which is a core feature of hyperontologies.

Describing experimental processes in the biomedical domain also requires a plurality of independent interoperable ontologies. For instance, the Ontology for Biomedical Investigations (OBI) aims to provide a logic conceptual framework for describing biomedical investigations. This task involves interoperability across several ontologies. For example, describing a PCR process involves at least OBI and ChEBI: *buffer*, *reagent* and *phenol* from ChEBI; *thermal_cycler*, *temperature_control_bath* and others from OBI. These classes are usually brought together via either OWL imports of the full ontologies (leading to a size explosion and the accompanying decrease in performance for reasoning) or simply by “slicing” the ontologies and putting together the classes on a need-to-have basis according to the MIREOT methodology [2]. This mechanism is facilitated by tools such as OntoFox [22], which allows users to input terms, fetch selected properties, annotations, and certain classes of related terms from source ontologies and save the results using the RDF/XML serialization of OWL. These hand-selected modules of external ontologies are then brought manually into the target ontology through imports, and the procedure has to be repeated every time the source ontology changes. The hyperontology framework can complement this by a transparent

and automated mechanism to achieve the required interlinked modules, and importantly, extracts modules based on logical principles rather than user steered “slicing”. It remains to be explored how these approaches can be combined and benefit from each other.

5 Outlook

We have sketched the Hyperontology framework and its sophisticated heterogeneous structuring mechanisms, and tried to illustrate their applicability to the domain of Biomedical Ontology by discussing several modelling scenarios in which heterogeneity is a central concern.

Biomedical ontologies, with their complex and heterogeneously interlinked sources of data, conceptual, spatial, and other kinds of knOWLedge, probably comprise the most complex application field for ontology engineering today.

In practice, biomedical ontologies often rely on simple subsumption hierarchies (**is_a**), aiming for a generic set of terms and their relationships. Even this, however, requires a strong methodological guidance as most biomedical ontologies are being developed by distributed groups, where teams do not necessarily follow the same classification systems. Such diversity makes a straightforward integration difficult, in particular when moving on to more highly axiomatised ontologies. Biomedical ontologies could evolve and grow in the way they did partly because the use of these controlled vocabularies has mostly been for annotation purposes. In addition, these ontologies are rich in lexical definitions, but not so much in terms of logical axioms. Lexical definitions are surely very important, they make it easier for ontology engineers to elicit knowledge and to manage and understand complex domains. However, a main purpose of formal ontologies is to represent human knowledge so that computers can interpret and reason with it, and not just to facilitate communication across human agents.

When acquiring new knowledge that is beyond the expressivity of is a hierarchies, interoperability across ontologies demands the network of ontologies to rely on a corresponding ‘network of axioms’. Reusing terms from one ontology in a new context cannot rely on simple

‘slicing’ or ‘cutting out’ operations, but has to be based on a logically well-founded method of ‘connecting’ the respective terms. The proposed use of the hyper-ontology framework is intended to enable just such a federated, interoperable solution, focusing on the logical definition of terms and a systematic and structured re-usability of axiomatisations. We are aware that we have only sketched some initial contributions towards this goal; much has to be done to make this really work and get feedback how well it scales in practice.

Acknowledgements

Work on this paper has been supported by the DFG-funded collaborative research centre SFB/TR 8 ‘Spatial Cognition’ and by the German Federal Ministry of Education and Research (Project 01 IW 07002 FormalSafe).

References

1. Baker, C. J. O., Kanagasabai, R., Ang, W. T., *et al.* Towards ontology-driven navigation of the lipid bibliosphere. *BMC Bioinformatics* 9 (2008).
2. Courtot, M., Gibson, F., Lister, A. L., Malone, J., Schober, D., Brinkman, R. R., and Ruttenberg, A. MIREOT: The minimum information to reference an external ontology term. *Applied Ontology* 6 (2011), 23–33.
3. de Matos, P., Alcántara, R., Dekker, *et al.* Chemical Entities of Biological Interest: an update. *Nucl. Acids Res.* 38 (2010), D249–D254.
4. Faratian, D., Clyde, R. G., Crawford, J. W., and Harrison, D. J. Systems pathology – taking molecular pathology into a new dimension. *Nat Rev Clin Oncol* 6, 8 (Aug 2009), 455–64.
5. Franke, M., Klein, P., and Schröder, L. Ontological semantics of standards and plm repositories in the product development phase. In *Global Product Development. Proc. 20th CIRP Design Conference 2010* (2010), A. Bernard, Ed., Springer, pp. 473–484.
6. Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., and Schneider, L. Sweetening Ontologies with dolce. In *Proc. of EKAW 2002* (2002), vol. 2473 of LNCS, Springer, pp. 166–181.
7. Golbreich, C., Horridge, M., Horrocks, I., Motik, B., and Shearer, R. OBO and OWL: Leveraging Semantic Web Technologies for the Life Sciences. In *Proc. of the 6th ISWC* (2007), vol. 4825 of LNCS, Springer, pp. 169–182.
8. Hastings, J., Dumontier, M., Hull, D. *et al.* Representing chemicals using OWL, description

- graphs and rules. In *Proc. of OWLED* (2010).
9. Hoehndorf, R., Batchelor, C., Bittner, T., *et al.* The RNA Ontology (RNAO): An ontology for integrating RNA sequence and structural data. *Applied Ontology* 6, 1 (2011), 53–89.
 10. Kalyanpur, A., Parsia, B., Horridge, M., and Sirin, E. Finding all Justifications of OWL DL Entailments. In *Proc. of ISWC/ASWC* (2007), vol. 4825 of *LNCS*, Springer, pp. 267–280.
 11. Kohlhase, M., Lemburg, J., Schröder, L., and Schulz, E. Formal management of cad/cam processes. pp. 223–238.
 12. Kutz, O., Mossakowski, T., and Lücke, D. Carnap, Goguen, and the Hyperontologies: Logical Pluralism and Heterogeneous Structuring in Ontology Design. *Logica Universalis* 4, 2 (2010), 255–333.
 13. Le Novère, N., Courtot, M., and Laibe, C. Adding semantics in kinetics models of biochemical pathways. In *Proceedings of the 2nd International Symposium on experimental standard conditions of enzyme characterizations* (2007).
 14. Mossakowski, T., and Kutz, O. The Onto-Logical Translation Graph. Tech. rep., University of Bremen, 2011. Submitted.
 15. Mungall, C. J., Bada, M., Berardini, T. Z., Deegan, J., Ireland, A., Harris, M. A., Hill, D. P., and Lomax, J. Cross-Product Extensions of the Gene Ontology. *Journal of biomedical informatics* (Feb. 2010).
 16. Rubin, D. L., Shah, N. H., and Noy, N. F. Biomedical ontologies: a functional perspective. *Briefings in Bioinformatics* 9, 1 (2008), 75–90.
 17. Sioutos, N., de Coronado, S., Haber, M. W., Hartel, F. W., Shaiu, W.-L., and Wright, L. W. NCI Thesaurus: A semantic model integrating cancer-related clinical and molecular information. *Journal of Biomedical Informatics* 40, 1 (2007), 30–43.
 18. Smith, M. K., Welty, C., and McGuinness, D. L. *The Web Ontology Language*, 2010.
 19. The Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nat. Genet.* 25 (2000), 25–9.
 20. The Gene Ontology Consortium. *The OBO language*, version 1.2, 2010.
 21. Villanueva-Rosales, N., and Dumontier, M. Describing chemical functional groups in OWL-DL for the classification of chemical compounds. In *Proc. of OWL: Experiences and Directions (OWLED 2007)* (2007).
 22. Xiang, Z., Courtot, M., Brinkman, R. R., Ruttenger, A., and He, Y. Ontofox: web-based support for ontology reuse. *BMC Res Notes* 3 (2010), 175.