

An Architecture for Semantic Navigation and Reasoning with Patient Data - Experiences of the Health-e-Child Project

Tamás Hauer¹, Dmitry Rogulin¹, Sonja Zillner², Andrew Branson¹,
Jetendr Shamdasani¹, Alexey Tsymbal², Martin Huber²,
Tony Solomonides¹, and Richard McClatchey¹

¹ CCS Research Centre, CEMS Faculty, University of the West of England
Coldharbour Lane, Frenchay, Bristol BS16 1QY, UK

² Corporate Technology Division. Siemens AG, Germany
first.last@{cern.ch,siemens.com,uwe.ac.uk}

Abstract. Medical ontologies have become the standard means of recording and accessing conceptualized biological and medical knowledge. The expressivity of these ontologies goes from simple concept lists through taxonomies to formal logical theories. In the context of patient information, their application is primarily annotation of medical (instance) data. To exploit higher expressivity, we propose an architecture which allows for reasoning on patient data using OWL DL ontologies. The implementation is carried out as part of the Health-e-Child platform prototype. We discuss the use case where ontologies establish a hierarchical classification of patients which in turn is used to aid the visualization of patient data. We briefly discuss the treemap-based patient viewer which has been evaluated in the Health-e-Child project.

1 Introduction

Digitized information management has greatly improved clinical practice during the past decades. Much patient data from demographic information to lab results to diagnostic images is now being stored in computerized form. Today, one of the main challenges for clinical information systems is to find, select and present the right information to the clinician from the vast amount of data that is available. This is a daunting task unless effective filtering, classification and visual aids are available. In this paper we consider an architecture for semantic navigation and reasoning with patient data, and share our experiences obtained within the EU FP6 project *Health-e-Child*¹. The functionality of the architecture hinges on the patient data stored in a database distributed over the Grid, a domain ontology with the knowledge relevant for lexicographic classification of patients, and two key data analysis components, for ontology-based reasoning and visualization. The focus in this paper is on presenting the later two; reasoning with ontologies

¹ <http://www.health-e-child.org>

and ontology-based visualization, which form the backbone of the considered architecture.

There has been lately much work on *ontology visualization* [1,2,3,4] that helps the user display and navigate underlying ontological concepts, see [5] for an extensive survey. In contrast to the mainstream works in the area, the present work proposes not the navigation of the ontology directly, but rather the visualization of instance data with the help of the knowledge that is represented by the ontology, or the deduced knowledge, projecting the respective data onto the ontology of interest. We employ two techniques suitable to ontology-based visualization of projected instance data for that; *facet browsing* and *treemaps* [6]. Each technique has its own benefits and limitations, which complement each other for the two techniques.

Visualization is tightly coupled with the *reasoning* component. Reasoning with ontologies is currently under active study in the semantic web field [7]; with biomedicine being one of the most popular application domains [8,9]. The ability to reason, that is to draw inferences from the existing knowledge to derive new knowledge is an important element for modern systems based on ontologies [7]. In particular, as we demonstrate in this paper using DL reasoning, reasoning with ontologies can help establish a hierarchical classification of patients for their intuitive visualization. By aligning patient data with relevant (fragments of) ontologies and inferring more descriptive patient ontology, improved patient data visualization and comparison can be realized.

The work in our study has been performed as part of the Health-e-Child (HeC) project. HeC is an EU-funded Framework Programme 6 (FP6) project, which was started in 2006, and aims at improving personalized healthcare in selected areas of paediatrics, particularly focusing on integrating medical data across disciplines, modalities, and vertical levels such as molecular, organ, individual and population. The project of 14 academic, industry, and clinical partners aims at developing an integrated healthcare platform for European paediatrics while focusing on some carefully selected representative diseases in three different categories; paediatric heart diseases, inflammatory diseases and brain tumours. The material presented in this paper contributes to the development of decision support facilities within the platform prototype which provide the clinicians with a tool to easily retrieve and navigate patient information and help visualizing interesting patterns and dependencies that may lead, besides personalized decision support concerning appropriate treatment, to establishing new clinical hypotheses and ultimately discovering novel important knowledge.

The paper is organized as follows. In section 2, we look at the requirements we have elicited from collaborating clinicians. In section 3, we review a few medical ontologies which we found useful for our use-cases. Section 4 contains the technical details of our approach of integrating patient data with external knowledge, and section 5 presents the architecture of our prototype platform. In Section 6, we visualize the hierarchical classification using treemap views.

2 Visualization Requirements in Clinical Practice

Clearly arranged visualization of patient data, supports clinicians in their daily tasks of clinical care and medical research. In order to navigate, analyze and visualize the dataset, it is useful to structure information and impose automatic annotation of patient records. In the following sections, we will illustrate how we realized the ontology-based visualization of patient information establishing the backbone of the introduced architecture. We will first describe the particular requirements for visualization of patient data in clinical practice. Bearing in mind the clinical visualization requirements, we will then sketch how we selected relevant medical knowledge sources and how treemaps in combination with our inferred patient ontology can be used for discovering correlations in patient records.

From extensive discussions with clinicians collaborating in Health-e-Child, we have learnt that the clearly arranged presentation of similar patients with respect to the complex and heterogeneous patient data becomes a key requirement for clinical decision support systems. For clinicians, the comparison of similar patients is a valuable source of information in the process of decision making. Therefore clinicians and medical researchers show a particular interest in visual means for comparing and analyzing the heterogeneous data of similar patients that cover demographics, family history, lab results, echocardiograms, MRI, CT, angiograms, ECG, genomic and proteomic samples, history of appointments, and treatments. Our existing data captures for each patient record more than 100 attributes describing the patient history and status data and allows the clinician to analyze patient records at a time. Our requirements elicitation has revealed the following further requirements in aiming for improved patient data visualization:

1. The discovery of patterns and dependencies in patient data. For example, establishing a correlation between the attributes “quality of life” and “tumour location” of brain cancer patients, is a routine task for clinicians. Therefore, the visualization of correlations between selected patient attributes becomes crucial in the clinical decision making process. As patient data attributes are provided in different levels of detail and precision, e.g. tumour location can be specified as “Cerebral Hemisphere” or, more detailed, as “Frontal Left Cerebral Hemisphere”, the computation and visualization of data attributes for correlation needs to reflect the variety of detail and precision.
2. The comparison of similar patient records with regard to relevant patient attributes should be supported by browsing facilities over the set of all patient records along context relevant features. Again, the browsing facilities need to reflect the variety of data in detail and precision.

Within traditional applications, users may browse and visualize patient data but little or no help is given when it comes to interpretation because the required semantics are implicit and thus inaccessible to the system. Hence, aiming for the means to enable the easy browsing of patient data and the visualization of complex information e.g. correlations for the establishment of new hypotheses,

we are integrating medical ontological knowledge to align patient data with the imposed knowledge structure thereby inferring the correct classification patient records.

3 Identified Medical Ontologies

For a beneficial integration of external semantics, one has to decide which external knowledge sources are appropriate for the purpose in mind, i.e. which external knowledge source captures relevant and helpful knowledge for a particular context. In our case, we aim to provide experts in brain tumour diagnosis and treatment with an improved method for patient data visualization and comparison. More precisely, we aim for the classification of patients with respect to their diagnosed tumour location or WHO tumour classification. We have chosen to use the Foundational Model of Anatomy (FMA) [10] covering the partitive hierarchy of brain regions as relevant and valuable medical knowledge. The coverage of the FMA is very comprehensive, containing approximately 70 thousand distinct anatomical concepts with more than 1.5 million relationships of 170 relationship types. We rely on fragments covering the concepts and relationships relevant to a particular visualization use-case. In our scenario, the established fragment encompasses all anatomical concepts describing possible brain tumour locations hierarchically structured by the `regional_part_of` relationship.

As second medical knowledge source, we identified the WHO classification of Tumours of the Nervous System establishing a classification and grading of human tumours that is accepted and used worldwide [11]. Its constituted entities establish a hierarchical structuring of histological typed tumours covering a multiplicity of factors, such as the immunohisto-chemistry aspects, genetic profiles, epidemiology, clinical signs and symptoms, imaging, prognosis, and predictive factors. The WHO classification of tumours refers to the ICD-O (International Classification of diseases for oncology) code and includes a WHO-grading scheme that is used for predicting response to therapy and outcome. For improved patient data visualization, we revert to its hierarchical structuring and its grading scheme. Similar to the integration of FMA's knowledge about tumour locations, we use the WHO classification's inherent hierarchical structuring for hierarchically representing patients' data.

4 Patient Record Classification by Reasoning

The Health-e-Child demonstrator is an integrated system that is built on top of a distributed platform, powered by grid technology, which hosts a distributed database and encompasses high-level enabling applications that exploit the intergrated medical data. The medical data is stored according to the Health-e-Child integrated data model and a service-oriented architecture provides access to storage, management and querying of the hosted information.

	Sex	Age at Diagnosis	...	Tumour Site	...
patient1	M	5	...	Cerebellum	...
patient2	F	7	...	Hypothalamus	...
...

Fig. 1. Simple view of patient data

For the richest possible interpretation, part of the medical data is annotated [12] using selected medical ontologies which allows for integration with external information and in some cases reasoning with expressive ontologies. In particular, when the external knowledge is expressed in some description logic then we can make use of available reasoners to provide enhanced, semantic query answering.

Our example use case is related to the visualization of the hierarchical classification of tumour patients. (Similar use cases have been discussed in [13,6].) We start with a simplified view of the patient database in Fig. 1; the HeC integrated data model is, of course, much more complex but the creation of such simple views is trivial. This database view is “flat” in that all attributes are individual (discrete) labels on the patient without explicitly defined semantics or structuring. Some of the attributes are numeric, others are categorical, taking values in a finite set of predefined concepts, for example tumour location, which refers to an anatomical region of the human body. In order to get correct answers to queries like “does patient x have a tumour in the Hindbrain”, the system must have access to and be able to reason with the knowledge about the partitive anatomy of the brain.

The Foundational Model of Anatomy in OWL. The tumour site in the HeC database is annotated [12] with concepts from the FMA which duly encodes the meronymy of human anatomical structures. The FMA is an originally frame-based ontology but there has been effort to convert it to OWL [14,15]. DL Reasoning, unfortunately, does not scale well with the size of ontologies thus in practice one has to identify manageable fragments that are suited both to the use-case and to the data at hand. We have experimented with locality-based fragments [16,17] which are natural choice for reasoning tasks as they guarantee logical consistency but occasionally result in too large fragments and the algorithm is slow. Graph-based fragments (e.g. [18]) have the advantage of performance in terms of speed and fragment size but there is no guarantee about logical completeness. We acknowledge that the selection of such fragments to use (semi-automatic or manual) and the segmentation of the ontology is a difficult task on its own; for the purpose of the description of the architecture in this paper, however, we assume that such fragments are already available. Back to our example, we make use of the `regional_part_of` subontology rooted at the concept `Brain` of [14]². The first ingredient to our use-case is the *FMA T-box*:

² [14] is a stratified ontology with an OWL DL and an OWL full part. Although the `regional_part_of` semantics is encoded in the OWL full part, the fragment we need is purely DL.




Fig. 2. Brain anatomy as defined in the FMA

$$\begin{aligned}
 & \text{Tr}(\text{regional_part_of}) \quad (\textit{Transitive}) \\
 & \geq 1 \text{ regional_part_of} \sqsubseteq \text{Anatomical_Structure} \quad (\textit{Domain}) \\
 & \top \sqsubseteq \forall \text{regional_part_of. Anatomical_Structure} \quad (\textit{Range}) \\
 \text{Cerebellum} & \sqsubseteq \text{Anatomical_Structure} \sqcap \exists \text{regional_part_of. Metencephalon} \\
 \text{Metencephalon} & \sqsubseteq \text{Anatomical_Structure} \sqcap \exists \text{regional_part_of. Hindbrain} \\
 \text{Hindbrain} & \sqsubseteq \text{Anatomical_Structure} \sqcap \exists \text{regional_part_of. Brain} \\
 & \vdots
 \end{aligned}$$

This ontology is our external knowledge that is independent of the Health-e-Child system. In order to make use of it, the information in our database has to be aligned with this ontology. The alignment is in part provided by the *Patient T-box*:

$$\begin{aligned}
 & \geq 1 \text{ has_tumour_location} \sqsubseteq \text{Patient} \quad (\textit{Domain}) \\
 & \top \sqsubseteq \forall \text{has_tumour_location. Anatomical_Structure} \quad (\textit{Range})
 \end{aligned}$$

This terminology establishes the semantics for the database records (*Patients*) and provides the alignment with the external ontology through the range of the *has_tumour_location* property. Finally, the relational instance data in Fig. 1 has to be mapped to conformant DL syntax, constituting the *Patient A-box*:

$$\begin{aligned}
 & \langle \text{patient1}, \text{Cerebellum}, \dots \rangle \\
 & \langle \text{patient2}, \text{Hypothalamus}, \dots \rangle \\
 & \quad \dots \\
 & \quad \Downarrow \\
 & (\text{Patient} \sqcap \exists \text{has_tumour_location. Cerebellum})(\text{patient1}) \\
 & (\text{Patient} \sqcap \exists \text{has_tumour_location. Hypothalamus})(\text{patient2}) \\
 & \quad \dots
 \end{aligned}$$

The queries for our use case are subclasses of the *Patient* class. When establishing the hierarchical classification of patients based on tumour location, these queries make up our *classification T-box* whose defined concepts are the labels for the patient classes:

$$\begin{aligned} \text{BrainTumourPatient} &\equiv \text{Patient} \sqcap \exists \text{has_tumour_location}. \\ &\quad (\text{Brain} \sqcup \exists \text{regional_part_of}.\text{Brain}) \\ \text{CerebellumTumourPatient} &\equiv \text{Patient} \sqcap \exists \text{has_tumour_location}. \\ &\quad (\text{Cerebellum} \sqcup \exists \text{regional_part_of}.\text{Cerebellum}) \\ &\vdots \end{aligned}$$

In effect, reasoning with the anatomy ontology ensures that patients will be classified/annotated as e.g. *cerebellum tumour patients* for every case that had originally been annotated with cerebellum or with any known regional part of the cerebellum as the tumour site.

To summarize, we consider the information at hand be represented in terms of DL, where we isolate three parts of the T-Box, a Patient terminology, which gives immediate semantics to the instances, an external knowledge base which is used as annotation and a classification ontology which is simply a collection of meaningful queries and is dictated by the use-case. In turn we arrive at the appropriate semantic classification of patients.

5 Prototype Platform Architecture

Let us now turn to our proposed architecture and implementation. The three core pillars of turning the above theory into a working infrastructure are: an ontology manager that can integrate the different knowledge components; a mapping mechanism between the database schema and the OWL DL A-box; and a reasoner which answers DL queries. These are implemented as services in the HeC platform service layer. Furthermore, these interact with additional components which are dictated by the use case, like a visualization component in our example.

Our implementation is based on the OWL API³ and some of the design decisions have been influenced by the conformance to it:

- The *OWL-DL Integrator* is a generalized `OWLontologyManager` which is responsible for importing and managing all the ontology components and loading the knowledge into a reasoner.
- The *DB/OWL DL Mapping component* creates simple views like that of Fig. 1 from the database and maps them to OWL-DL.
- The *Reasoner* uses the set of assertions and knowledge accumulated above and answers semantic queries, in particular creates the inferred patient classification.
- An optional *Ontology Transformation Component* is used in the hierarchical classification use-case to establish the set of queries or labels (e.g. “Cerebellum Tumour Patients”).
- The *Interpretation & Visualization Component* maps the inferred classification from OWL to the appropriate representation of the user interface.

³ <http://owlapi.sourceforge.net/>

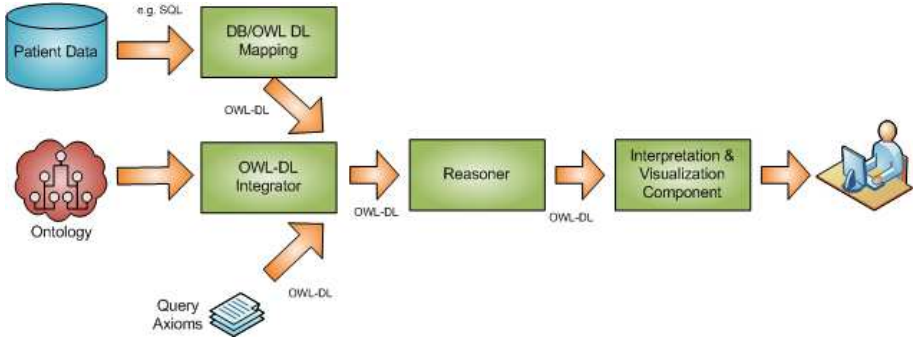


Fig. 3. Prototype Architecture

We briefly discuss these components in turn. The *DB/OWL DL Mapping component* uses the semantic annotations of the patients’ data to expose patients’ information as OWL DL ontology. First, a flat view is created from the relevant relations which includes the entity identifiers (patient ID), the concept URIs for the hierarchical classification (tumour location) and additional attributes (e.g. status at the end of treatment in later examples) which are of interest but don’t contribute to the reasoning. Second, the relevant columns of the relation are translated into DL which is expressed in OWL. This is implemented using a Mapper class which is governed by the Patient terminology and a set of mapping descriptions which bridge the relational and DL schema. When browsing along multiple axes is required, they are all included in the OWL view.

The *OWL-DL Integrator* is a generalized *OWL Ontology Manager*. It implements (exposes) multiple ways of accumulating knowledge, including loading OWL from external URI, loading instance data from the database using Mapper instances and adding standalone axioms on the fly. It populates the reasoner with the merged external, patient and classification ontologies and initializes the reasoning.

The *Reasoner* creates the inferred hierarchical classification of patients. We currently use the Pellet 1.5 reasoner engine in the Health-e-Child platform prototype.

The transitive `regional_part_of` property on anatomical concepts induces the subsumption relationship on patient classes (`has_t1 ≡ has_Tumour_Location`):

$$\begin{aligned}
 X &\sqsubseteq \exists \text{regional_part_of}.Y \\
 &\Downarrow \\
 \exists \text{has_t1}.\exists \text{regional_part_of}.X &\sqsubseteq \exists \text{has_t1}.\exists \text{regional_part_of}.Y
 \end{aligned}$$

We can’t, however, directly exploit this inference because the classes are not defined: the reasoner only creates inferred subsumption hierarchy for named classes and the visualization also requires the definition of the classes for the lexicographic hierarchy. In other words, we need to supply the set of queries that govern the classification.

To that end, the *ontology transformation component* creates the missing bit: the definition of class names for the hierarchy. In our example this is a set of name definitions based on the corresponding concept names. Creating this set is trivial because we obviously don't use any structure of the ontology.

The *Interpretation & Visualization Component* transforms the inferred OWL into the format conformant to the API of the visualization software. It can also add further attributes from the database which were not part of the reasoning process.

6 Patient Data Visualization Using TreeMaps

The user requirements for visualization patient data revealed the importance of means for discovering correlations in patient data. Therefore, we decided to use treemaps as visualization component. Treemaps [19] are an efficient two-dimensional technique to visualize hierarchical data structures. It is particularly popular for disc storage view of hierarchical filesystems because file size is an aggregate attribute of files. It is a space filling technique, i.e. one that uses the entire screen area by dividing it up between leaf nodes which are subsequently grouped into enclosing rectangles [5]. The image is effectively a rectangular Venn-diagram of nonintersecting sets. Besides the set semantics, there are other attributes such as colour, choice of font and label which can represent attributes of the data beyond that of the hierarchy.

To allow for improved patient visualization and discovery of patient correlation, we represent patients as rectangles of equal size so that the cardinality of patient classes can be easily seen. Figure 4 shows the user interface based on the patient taxonomy that has been inferred from the anatomical meronomy. The medical background information, in our example the hierarchical structured brain tumour locations establishes enclosing and nested areas that are labelled by regions of the brain. Colour may be used to visualize further attributes, in our example the status at the end of treatment for each patient. It is meant to show that space-filling hierarchies can be useful for visualizing correlation, in this example one finds that patients with tumour in the Hindbrain tend to have better prognosis than in other areas (especially the Left Temporal Lobe). In a similar manner, we can correlate and visualize patient record with respect to the "WHO tumour classification" and, for instance their "quality of life" parameter.

A further advantage of treemaps is capability for "binning" the data. An example of Treemap with patient binning would show the patients grouped not only by tumour location but by age as well. It follows visually in Fig. 5 that – in this sample – the end of treatment status of younger patients is worse than that of older patients: there is noticeable correlation which the user can discover by visualizing the data the right way.

While ontology-based visualization using treemaps was already reported, for example to visualize and navigate the Gene Ontology and microarray data [20], we are unaware of its applications in clinical decision support, where patient proximity is visualized by projecting patient data onto existing ontologies.

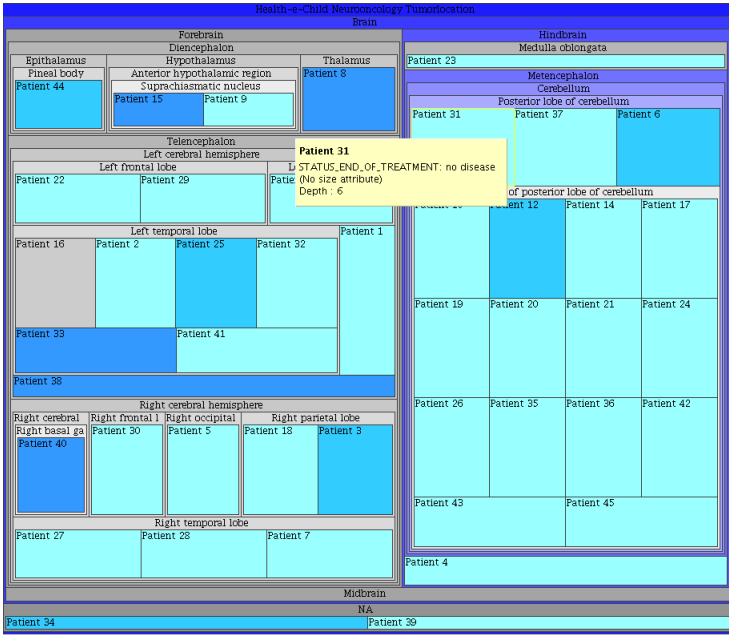


Fig. 4. A full Treemap view of the ontology assisted data representation

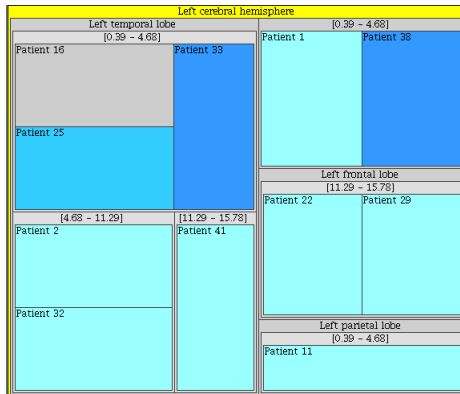


Fig. 5. Binning: patients are grouped both by tumour site as well as age

6.1 Related Approaches in Knowledge Visualization

The work presented in this section relates to visualization for knowledge discovery. The motivating example is that of the physician who believes that an extensive set of patient data would reveal subtle patterns if only it could be visualized in the right way. The right way is here taken to mean analyzed in

terms of established or even tentative concepts. Thus the ontology carries what may have already been accepted or adduced as a research hypothesis, and the data is visualized on that basis, as a means of strengthening the evidence or as a means of bolstering the new hypothesis.

There is a significant volume of work on visualization with (and without) ontologies. It may be differentiated from the work presented here on the basis of various criteria, such as: what is visualized, why and how. Visualization may be of data from a homogeneous database in order to display relative volume or from a heterogeneous source to expose some other feature or for exploratory data analysis [21]; it may be of the structure of the data (rather than the data itself) in order to reveal class-level relationships [3]. Visualization has been used to facilitate query formulation or to order threads of data in some schematic way (e.g. temporally); to display a data schema or to inform navigation through the data [22]. In particular cases, ontology-based visualization has been used to support queries based on temporal abstractions [23]; to enrich maps with additional geographic information [24]; to reveal multiple levels of abstraction in decision-tree generation [25] and to assist in information mining [26]; very popularly, to map social networks and communities of common interest [27], [28], [29]. Ontologies have also been used for knowledge discovery without visualization, especially in the integration of heterogeneous scientific repositories ([30]).

For an up to date survey of ontology visualization methods reference must be made to [5]. This paper provides a near exhaustive discussion of methods of classification, of visualization techniques (Euclidean, hyperbolic and spherical space, node-and-link, and other less obviously geometric methods), of representation, overview and focus methods, but it is distinctive in offering no discussion of content. There is a body of work dating back to the mid-'90s on database visualization (see [21]) some of it devised to assist with data mining tasks.

Arguably the work closest to ours in spirit is [31], although their approach is designed more to manage the heterogeneity of web sources in order to summarize the information provided. Their visual method is based on node-and-link representation and resembles mind-mapping. On the other hand, [32] use 'semantic visualization' to support the development process for database-oriented systems, but in the process must tackle questions of data visualization.

7 Conclusions and Future Work

As the biomedical and clinical information available in digitized form is continually increasing, so is the demand for advanced data integration and visualization tools to help clinicians and researchers to explore the information and knowledge at their hands. In this paper we presented a system for patient data visualization which combines reasoning about data using external knowledge with advanced visualization that makes use of the inferred structure. Our approach makes use of existing ontology visualization paradigms with a view to presenting the data whose semantics is governed by the ontology.

We have presented an architecture of a system implementing our framework and demonstrated its functionality using paediatric brain cancer data acquired in Health-e-Child together with anatomical knowledge from the Foundational Model of Anatomy for the purpose of creating annotated treemap views of patient characteristics.

We presented a demonstrator of this framework to clinicians of the Health-e-Child project to evaluate early user experience. We demonstrated the use case of treemap visualization presented in this paper and facet browsing of [6]. In both cases the response was positive: facet browsing helps the clinicians to locate follow-up patients based on incomplete information while the treemap was appreciated for the easy visualization of correlation between clinical attributes. The clinicians also noted that it has potential in education and training.

There are several challenges along the lines we presented. A number of approaches exist for visualizing hierarchies but to really exploit their power, end-users have to be trained so that they can define the visualization which suits their interest and then navigate it. Inferring the hierarchy of the data based on its semantics requires suitable ontologies and alignment with them. In the medical domain, annotation with the UMLS meta-thesaurus is a good choice because it maps to many ontologies. A serious bottleneck is the reasoner as performance scales very badly with the size of the ontologies. We compensated this with fragment extraction but ontology segmentation is also difficult and, especially when one aims for consistent fragment, it is slow and can only be done off-line.

Our future work includes a comparative analysis of how different ontology fragments perform in this approach with real data and we are investigating how this classification scales with the number of patients. We also want to explore the means to give some control over the selection of ontology fragments and definition of classification criteria to the end-users. This is a difficult problem but could prove very interesting.

Acknowledgements

This work has been partially funded by the EU project Health-e-Child (IST 2004-027749). The authors wish to acknowledge support provided by all the members of the Health-e-Child consortium in the preparation of this paper.

References

1. Bosca, A., Bonio, D.: Ontosphere: more than a 3d ontology visualization tool. In: SWAP The 2nd Italian Semantic Web Workshop (2005)
2. Noy, N.F., Fergerson, R.W., Musen, M.A.: The knowledge model of protege-2000: Combining interoperability and flexibility. In: Dieng, R., Corby, O. (eds.) EKAW 2000. LNCS (LNAI), vol. 1937. Springer, Heidelberg (2000)
3. Mutton, P., Golbeck, J.: Visualization of semantic metadata and ontologies. In: IV 2003. Proceedings of the Seventh International Conference on Information Visualization, Washington, DC, USA, p. 300. IEEE Computer Society, Los Alamitos (2003)

4. Pietriga, E.: Isaviz, a visual environment for browsing and authoring rdf models. In: Eleventh International World Wide Web Conference Developer's Day (2002)
5. Katifori, A., et al.: Ontology visualization methods a survey. In: *ACM Comput. Surv.*, vol. 39, p. 10. ACM, New York (2007)
6. Zillner, S., Hauer, T., Rogulin, D., Tsymbal, A., Huber, M., Solomonides, T.: Semantic Visualization of Patient Information. In: *Computer Based Medical Systems (CBMS)* (2008)
7. van Harmelen, F., Huang, Z., ten Teije, A.: Reasoning with inconsistent ontologies: framework, prototype, and experiment. In: Studer, R., Davies, J., Warren, P. (eds.) *Semantic Web Technologies*. Wiley, Chichester (2008)
8. Rubin, D.L., Dameron, O., Musen, M.A.: Use of description logic classification to reason about consequences of penetrating injuries. In: *American Medical Informatics Association Conference (AMIA 2005)*, pp. 649–653 (2005)
9. Dameron, O., Roques, E., Rubin, E., Marquet, G., Burgun, A.: Grading lung tumors using owl-dl based reasoning. In: *Proceedings of the 9th International Protege Conference* (2006)
10. Rosse, C., Mejino, J.: A reference ontology for bio-medical informatics: the foundational model of anatomy. In: *Biomedical Informatics*, vol. 36, pp. 478–500 (2003)
11. Louis, D., et al.: The 2007 WHO Classification of Tumours of the Central Nervous System. *Acta Neuropathol* 114, 97–109 (2007)
12. Rogulin, D., Ruiz, E.J., McClatchey, R., Berlanga, R., Hauer, T., Shamdasani, J., Zillner, S., Branson, A., Nebot, V., Manset, D., Freund, J.: Medical Data Integration and the Semantic Annotation of Medical Protocols. In: *Computer Based Medical Systems (CBMS)* (2008)
13. Tsymbal, A., Huber, M., Zillner, S., Hauer, T., Zhou, K.: Visualizing patient similarity in clinical decision support. In: Hinneburg, A. (ed.) *LWA*, pp. 304–311. Martin-Luther-University, Halle-Wittenberg (2007)
14. Noy, N.F., Rubin, D.L.: Translating the foundational model of anatomy into owl. Technical report, Stanford Center for Biomedical Research (2007)
15. Golbreich, C., Zhang, S., Bodenreider, O.: Migrating the FMA from Protege to OWL. In: *8th International Protege Conference*, Madrid (2005)
16. Grau, B.C., Horrocks, I., Kazakov, Y., Sattler, U.: Just the right amount: extracting modules from ontologies. In: *WWW 2007. Proceedings of the 16th international conference on World Wide Web*, pp. 717–726. ACM, New York (2007)
17. Sattler, U., Schneider, T., Jiménez-Ruiz, E., Grau, B.C., Berlanga, R.: Safe and economic re-use of ontologies: A logic-based methodology and tool support. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) *ESWC 2008*. LNCS, vol. 5021. Springer, Heidelberg (2008)
18. Jimenez-Ruiz, E., et al.: Ontopath: A language for retrieving ontology fragments. In: *OTM Conferences (1)*, pp. 897–914 (2007)
19. Shneiderman, B.: Tree visualization with tree-maps: 2d space-filling approach. *ACM Transactions on Graphs* 11, 92–99 (1992)
20. Babaria, K., Baehrecke, E.H., Dang, N., Shneiderman, B.: Visualization and analysis of microarray and gene ontology data with treemaps. *BMC Bioinformatics* 5(1), 84 (2004)
21. Keim, D.A.: Information visualization and visual data mining. *Transactions on Visualization and Computer Graphic* 8, 1–8 (2002)
22. Cannon, A., et al.: Ontology-driven automated generation of data entry interfaces to databases. In: *Key Technologies for Data Management (LCS)*. Springer, Heidelberg (2004)

23. Shahar, Y., Cheng, C.: Ontology-driven visualization of temporal abstractions. In: KAW 1998 (1998)
24. Ipfelkofer, F., et al.: Ontology Driven Visualisation of Maps with SVG - An Example for Semantic Programming. In: Conference on Information Visualization, pp. 424–429 (2007)
25. Zhang, J., et al.: Ontology-driven induction of decision trees at multiple levels of abstraction. LNCS (LNAI). Springer, Heidelberg (2002)
26. Castillo, J.A.R., et al.: Information extraction and integration from heterogeneous, distributed, autonomous information sources - a federated ontology-driven query-centric approach. In: Conference On Information Reuse and Integration, IEE International, pp. 183–191 (2003)
27. Mika, P.: Ontologies are us: A unified model of social networks and semantics. In: Gil, Y., Motta, E., Benjamins, V.R., Musen, M.A. (eds.) ISWC 2005. LNCS, vol. 3729, pp. 522–536. Springer, Heidelberg (2005)
28. Zheng, J., Niu, J.: Unified mapping of social networks into 3d space. In: Second International Multi-Symposiums on Computer and Computational Sciences (IM-SCCS 2007), pp. 305–311 (2007)
29. Domingue, J., Motta, E.: A knowledge-based news server supporting ontology-driven story enrichment and knowledge retrieval. In: Fensel, D., Studer, R. (eds.) EKAW 1999. LNCS (LNAI), vol. 1621, pp. 103–120. Springer, Heidelberg (1999)
30. Christophides, V., et al.: Ontology-driven integration of scientific repositories. In: 4th International Workshop on Next Generation Information Technologies and Systems, pp. 190–202. Springer, London (1999)
31. Fluit, C., et al.: Ontology-based information visualisation: Towards semantic web applications. In: Geroimenko, V. (ed.) Visualising the Semantic Web, 2nd edn. Springer, Heidelberg (2005)
32. Sheth, A., Avant, D.: Semantic visualization: Interfaces for exploring and exploiting ontology, knowledgebase, heterogeneous content and complex relationships. In: NASA Virtual Iron Bird Workshop (2004)