

Interpreting patient data using medical background knowledge

Heiner Oberkampff, Sonja Zillner, Bernhard Bauer, Matthias Hammon

Angaben zur Veröffentlichung / Publication details:

Oberkampff, Heiner, Sonja Zillner, Bernhard Bauer, and Matthias Hammon. 2012. "Interpreting patient data using medical background knowledge." In KR-MED: ICBO 2012 - Proceedings of the 3rd International Conference on Biomedical Ontology (ICBO 2012), Graz, Austria, 21-25 July 2012, edited by Ronald Cornet and Robert Stevens, 6. Aachen: CEUR-WS.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under the following conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publizieren/>



Interpreting Patient Data using Medical Background Knowledge

Heiner Oberkamp^{1*}, Sonja Zillner¹, Bernhard Bauer² and Matthias Hammon³

¹Corporate Technology Siemens AG, 81739 Munich, Germany

²Programming Distributed Systems, University of Augsburg, 86135 Augsburg, Germany

³University Hospital Erlangen, 91054 Erlangen, Germany

ABSTRACT

Clinical patient data, such as medical images and reports, establish the basis of the diagnostic process. In order to improve the access to heterogeneous and distributed clinical data sources recent work concentrates on extracting semantic annotations with links to concepts of medical ontologies, such as RadLex, FMA, SNOMED CT or others. However, these annotations are not always on the appropriate level of detail for clinical applications as they simply reflect the descriptive content of the respective patient data. For integrating the data into the clinical work-flow, an interpretation of annotations using medical background knowledge is needed. In this paper we present a use-case of such an interpretation: we have built an initial ontology containing lymphoma-related diseases and symptoms as well as *their relations*. The created ontology is used to infer likely diseases of patients based on annotations. In this way, annotations can be understood in the context of likely diseases and help the clinician to make a diagnosis. By means of a prototype implementation we evaluate our approach and identify further knowledge requirements for the model.

1 INTRODUCTION

Clinical patient data, such as medical images and patient reports, provide the basis for the diagnostic process. However, the enormous volume and complexity of data prevents clinical staff to get the full use of the content of the data by reviewing it all. Recent work aimed to make heterogeneous clinical data better accessible (for search or other processes) by means of structured semantic annotation with concepts of well established medical ontologies like RadLex, the FMA, SNOMED CT or others. For instance, the Theseus MEDICO project¹ aims at the automatic extraction of descriptive content of medical images for better integration into clinical processes as e.g. decision making. Approaches for semantic annotation range from automatic image parsing (Seifert *et al.*, 2009) and information extraction DICOM headers and structured reports (Möller *et al.*, 2009) to (aided) manual approaches (Wennerberg *et al.*, 2008), (Channin *et al.*, 2010) and (Rubin *et al.*, 2008). In the MEDICO project a dedicated Annotation-Schema, the MEDICO-Annotation-Ontology (Seifert *et al.*, 2010), was created to store the annotations of clinical data in a structured way in order to improve precision and recall in information retrieval in comparison to simple key-word tagging as pointed out in (Opitz *et al.*, 2010). Moreover, in MEDICO, annotations were used to classify lymphoma patients according to the Ann-Arbor staging system which relies mainly on the location of enlarged lymph nodes (Zillner, 2009).

So the good news is that great progress has been made concerning the extraction of annotations. The problem however is that these

annotations capture only descriptive information of its content, i.e. the observations made, the findings discovered, the various symptoms identified². Though in practice clinicians often want to search for higher level information, e.g. disease information. Consider the diagnostic process: after a CT-scan a clinician looks for findings indicating certain diseases he suspects the patient might have. For instance, the clinician looks for *cancer-indicating* findings or symptoms in the CT image series. The problem is that even though existing medical ontologies encompass information about diseases and symptoms, knowledge about their interrelations is missing.

Today methods for capturing and exploiting the information about semantic relation between symptoms, for instance CT-image annotations describing findings, and their associated diseases are missing. Thus, the search for *cancer-indicating* findings is only possible through a search for specific finding as e.g. "enlarged mediastinal lymph nodes" or "enlarged spleen" assuming that the clinician is informed about likely symptoms of a disease. However, clinicians are usually experts in one particular domain, leading to a lack of prior knowledge about the interrelations of symptoms and diseases in case certain diseases are no longer in the scope of their expertise. In other words, there is a clear danger that the information about the relevance of identified symptoms remains overlooked or misinterpreted, leading to wrong or not appropriate treatments, etc. Thus, the relevance-based highlighting of information about clinical observations in the context of likely diseases supports clinicians to improve their treatment decisions.

Within this paper, we introduce a use case scenario in the domain of clinical diagnose that relies on the seamless interpretation of annotations by integrating formalized medical background knowledge about interrelations of diseases and symptoms. However, the information about disease and symptom interrelations is not covered by existing medical ontologies. We fill the missing link by proposing a Disease-Symptom-Ontology that contains diseases, symptoms, *their relations* and as many links as possible to established medical ontologies. As a disease-symptom ontology can be arbitrarily complex, we need to be very specific about the scope and the goal of the proposed ontology. The contribution of this paper is to detail the knowledge model as well as the knowledge engineering steps of the first version of the ontology model that aims to fulfil the requirements of our use case scenario: the aim here is to infer likely diseases based on patient annotation data, helping the clinician in the diagnosis. After the automatic detection of an initial set of findings and corresponding diseases, this information can be used in differential diagnosis, where the clinician intends to either exclude or strengthen particular diseases. So the first list of symptoms and likely diseases in turn indicates to check for further symptoms. E.g. "enlarged lymph nodes" indicates to check for B-symptomatic i.e.

*To whom correspondence should be addressed: heiner.oberkampf.ext@siemens.com

¹ <http://theseus-programm.de/en/920.php>

² We exchangeably use the terms symptom, sign, finding and observation.

weight loss, fever and night sweat). Understanding available symptom annotations in the context of diseases improves the process significantly as the clinician will not miss any findings and symptoms of diseases that might be out of his (main) expertise. Through our application that relies on the Diseases-Symptom-Ontology we are able to make the likely diseases of the patient explicit, which again improves and accelerates the diagnostic process. Based on a clinical evaluation of the implemented use case, we identified further knowledge requirements for the model.

In Section 2, we will first sketch the related work of clinical decision support systems and medical ontologies as well as discuss in Section 3 the medical background-knowledge that is required for interpreting clinical annotations with respect to likely diseases. Section 4 shows how the medical knowledge can be represented in a Disease-Symptom-Ontology. By means of a prototype implementation we evaluate our approach and identify further knowledge requirements for the model (Section 5). We conclude with an outlook of the future work.

2 RELATED WORK

2.1 Clinical Diagnostic Support

Although various clinical diagnostic support systems have been realized (for a comprehensive overview we refer to the open-clinical website³), they are either tailored to a particular type or range of diseases, e.g. cardiac diseases (K and Singaraju, 2011) or aim to efficiently access and integrate to (only) external medical knowledge resources, such as the access to evidence-based best-practices⁴ or the access to relevant international studies and publications⁵. Moreover, various standalone decision support systems for clinical diagnosis, such as MYCIN (Shortliffe, 1976), CASNET (Kulikowski and Weiss, 1982), DXplain⁶ or other expert systems for diagnosis have been developed that require manual input of symptom and finding information to infer a set of likely diseases.

2.2 Medical Ontologies

We also analysed existing medical ontologies like SNOMED CT, FMA, RadLex and DOID with respect to disease-symptom relations. The Human Disease Ontology (DOID) consists of about 8000 diseases structured in a hierarchy, SNOMED CT provides a huge vocabulary of clinical terms, the Foundational Model of Anatomy (FMA) – a detailed description of the anatomy, RadLex – a structured terminology for the radiological domain ... However none of them contains diseases and symptoms *and* the relation between them, such that we could use the information to analyse our annotations with respect to diseases. E.g. SNOMED CT contains "Hodgkin-Lymphoma" and "lymphadenopathy" but no relation in-between. The only relation we have is that the "lymph node structure" is FindingSiteOf "Hodgkin-Lymphoma" as well as "lymph node structure" is FindingSiteOf several types of "lymphadenopathy". The Human Disease Ontology gives us a good hierarchy of

diseases and even though DOID aims also at ontological treatment of diseases and diagnosis (Scheuermann *et al.*, 2009) relations to symptoms, if present, are hidden in a owl:AnnotationProperty called "def" as text. For lymphoma there are no symptoms specified in DOID.

3 MEDICAL BACKGROUND KNOWLEDGE

Medical Ontologies provide a comprehensive and well structured vocabulary, which allows to describe precisely the content of images and reports in annotations. When we want to use these descriptive annotations (observations made, findings, symptoms) in clinical decision support systems, we need to interpret them. Existing medical ontologies however contain most of their information in their class-hierarchy. They do not have enough relations of symptoms to other relevant concepts as e.g. diseases, examinations or medication. In particular, existing ontologies can't be used to infer likely diseases, plan further treatments or control therapy results and medication effects based on annotations describing findings.

Similar to the cognitive decision process of clinicians, who rely on their experience and expertise, we take medical background knowledge to automatically interpret the findings and symptoms in order to integrate them in aforementioned clinical processes. In the first use case we aim to make annotations useful in diagnosis. Predominantly we need the relations between symptoms and diseases. This background knowledge is used to make likely diseases explicit (see figure 1). Knowing likely diseases makes it easier to plan next examinations. Further, the clinician gets an overview of what symptoms of certain disease are present without going through all images and reports. Especially this helps the clinician not to miss symptoms in the diagnostic process.

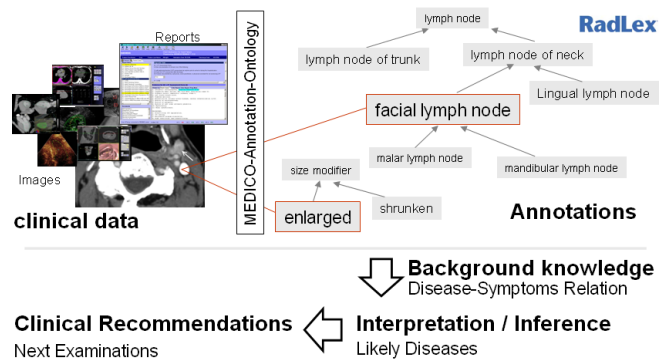


Fig. 1. From annotated patient data to diagnosis support: taking medical background knowledge to interpret annotations.

Knowledge about the disease-symptom-relations can be found in common clinical knowledge resources as for example "Innere Medizin" (G. Herold und Mitarbeiter), where about 300 diseases are listed and described. Besides definition of the disease with corresponding symptoms, epidemiological information, classification, risk-factors, clinical best practices for further treatment and therapy, the typical age and other useful information can be found in "Innere Medizin". Sometimes there is valuable extra-information as probabilities for occurrence of symptoms (e.g. "Hodgkin-Lymphoma" has the symptom "enlarged lymph nodes" with probability 85%) or

³ <http://www.openclinical.org/dss.html>

⁴ ZynxAmbulatory: <http://www.zynxhealth.com/Solutions/ZynxAmbulatory.aspx>

⁵ LeitsymptomNavigator: http://www.albis.de/home/news/nachricht-anzeigen/browse/3/datum////-1e3591bb94/?tx_ttnews%5BbackPid%5D=255&cHash=f625030d5cc0c3ad26d76d1a406e1945

⁶ <http://cs.mgh.harvard.edu/projects/dxplain.html>

hints for differential diagnosis (e.g. symptom "enlarged mediastinal lymph node" indicates differential diagnosis "Hilus-Tbc", "Non-Hodgkin-Lymphoma", "Bronchial carcinoma" and others). Some of the symptoms are precisely specified (e.g. "weight loss" more than 10% within the last 6 month).

4 THE DISEASE-SYMPATOM-ONTOLOGY

Our goal is to establish a knowledge model able to represent information contained in Herold's "Innere Medizin" necessary for clinical diagnosis. First of all this is the relationship between symptoms and diseases. As we want to use this Disease-Symptom-Model for the interpretation of annotations as shown in figure 1 we also choose to model the relations in an ontology with links to the medical ontologies which are used for annotations. Figure 2 shows how the Disease-Symptom-Ontology relates to other ontologies.




Fig. 2. Ontology architecture: the Disease-Symptom-Ontology has to be linked to all medical ontologies used for annotation of clinical patient data.

The cornerstones of the Disease-Symptom-Ontology are the classes DiSy:Disease and DiSy:Symptom with subclasses for specific diseases and symptoms. Additionally to the classes we have individuals for all diseases and symptoms to be able to relate symptoms and diseases with the owl:ObjectProperty DiSy:hasSymptom and a sub-property DiSy:hasLeadingSymptom. This approach is similar to the one in RadLex, the FMA or other medical ontologies which contain both classes and individuals. The relations should not be understood in the way "d always has symptom s" as this is not true in the medical domain. At this point we avoid getting into technical questions like how to represent formally that a disease "may" be caused by something or that a disease "may" show up by some symptom. For a good description of that problem we refer to (Rector *et al.*, 2008). Even though existing medical ontologies provide a detailed description of the medical domain this link is still missing as described in the related work section.

4.1 Establish Alignments

Linking the Disease-Symptom-Ontology to other medical ontologies has two reasons: firstly, this is necessary as we aim to interpret annotations which are taken from them. Our second purpose is to benefit from their structure (see structure import below). Figure 3 exemplarily illustrates how symptoms of DiSy are linked to RadLex such that annotations can be interpreted in a disease-context: a CT-image annotated with RadLex-concepts "enlarged" and "facial lymph node": we infer that the image shows a lymphoma-indicating finding but not a colorectal-cancer-indicating.

Links to other ontologies: Basically we hold the identifiers of the referenced concepts with owl:AnnotationProperty relations

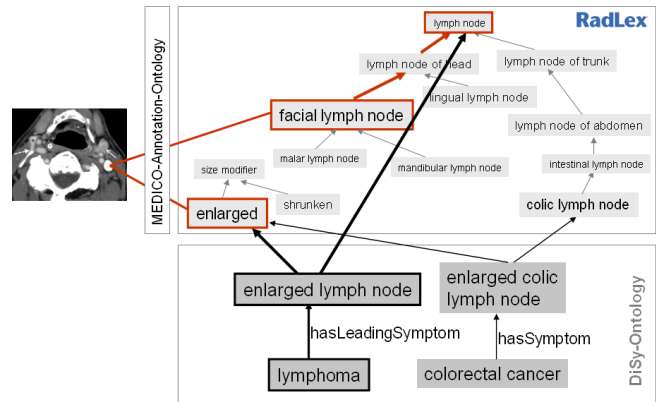


Fig. 3. A CT-image showing a lymphoma-indicating finding.

like DiSy:hasDOID, DiSy:hasRadLex_ID, DiSy:hasFMA_ID and DiSy:hasSNOMED_ID. This way of referencing works well if there exist corresponding concepts. However, finding corresponding concepts is generally difficult, so we take several approaches of linking concepts. We always try to link to single corresponding concepts but additionally we may link to two concepts, the one defining the location or parameter and the other – the modification: examples are "lymph node" and "enlarged" or "carcinoembryonic antigen" and "raised". If we do not find one single corresponding concept for a symptom, this is the only way of linking:

```
DiSy:Enlarged_mediastinal_lymph_node
  DiSy:hasSNOMED_ID SNOMEDCT:52324001;
  DiSy:hasAnatomicalRegion_RadLex RID28891;
  DiSy:hasModifier_RadLex RID5791.
```

Here SNOMED CT has a corresponding concept SNOMEDCT:52324001 ("mediastinal lymphadenopathy"), whereas RadLex has no single one, so we need to take the two concepts RID28891 ("mediastinal lymph node") and RID5791 ("enlarged") with respective relations. Using two links is similar to the annotation-techniques in the MEDICO project. Indeed, many different links to the ontologies allow us to recognize better the symptoms described in the MEDICO-annotations. The recognition is done through querying the annotation data in the following way: in the Disease-Symptom-Ontology we have a class DiSy:Patient and an owl:ObjectProperty :showsSymptom for relating patients with symptoms. Our query component searches for annotations, representing symptoms as shown in figure 3. If such annotation data is found, let's say for a symptom S, then a triple "patient :showsSymptom S" is added to our ontology. As we do not need direct matches of annotations and symptoms (as illustrated in figure 3) one can see that through the referenced ontologies we fan out our symptoms to all subclasses of the linked concepts: even though DiSy does not contain a symptom, which directly references "enlarged" and "facial lymph node", these image annotations will be captured at least as the symptom "enlarged lymph node", since "lymph node" is a superclass of "facial lymph node".

4.2 Importing Structure

As mentioned above we try to reuse knowledge from other medical ontologies. In this section we describe how the subclass-structure of DiSy is automatically created. All diseases are initially entered

as simple sub-classes of DiSy:Disease, i.e. without any hierarchical information. Likewise for symptoms. The only information we need to enter is the link to external ontologies. The big advantage of importing the hierarchy is that we reuse existing knowledge. If two referenced concepts in an external ontology (for diseases e.g. to DOID) stand in a sub-class relation, then we will enter a sub-class relation between the referencing concepts of DiSy. This can be achieved by a simple SPARQL CONSTRUCT statement. E.g. we have two diseases "Lymphoma" and "Hodgkin-Lymphoma", initially direct sub-classes of DiSy:Disease and with corresponding links to DOID. Due to DOID "Hodgkin-Lymphoma" is a sub-class of "Lymphoma", so this relation is added to our ontology. The subclass relation in the external ontology does not have to be direct, but it could also be a sub-class path. Symptoms are structured similarly (see figure 4): "Enlarged lymph node" and "enlarged mediastinal lymph node" - in RadLex we have a subclass path "mediastinal lymph node", "deep lymph node of thorax", "lymph node of thorax", "lymph node of trunk", "lymph node". As `rdfs:subClassOf` is a transitive relation after RDFS-reasoning within RadLex "mediastinal lymph node" becomes a subclass of "lymph node" so we make "enlarged mediastinal lymph node" a subclass of "enlarged lymph node". This is done again by SPARQL CONSTRUCT statements.

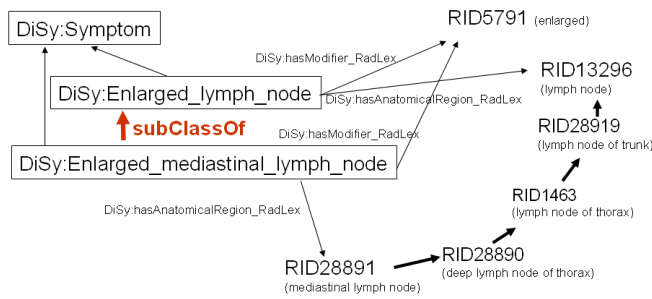


Fig. 4. Getting the subclass-hierarchy from established ontologies.

The structure import is done only once so that its computational effort can be accepted.

5 PROTOTYPICAL IMPLEMENTATION

5.1 Inference of likely Diseases

Given a patient with an initial set of symptoms (explicitly represented within the patient's annotations), we are aiming to infer a ranked list of likely diseases. From the Disease-Symptom-Ontology and the initial set of symptoms we can derive a list of likely diseases and in a second step for each disease a set of related symptoms. After aligning the annotation data with the set of related symptoms the set of related symptoms can be split into three categories for each disease:

- **Present symptoms:** symptoms for which corresponding annotations were found.
- **Absent symptoms:** symptoms which were under inspection but did not show up (e.g. "no enlarged lymph nodes in neck-area"). Absent symptoms are of high importance as they allow a clinician to exclude certain diseases.

- **Open symptoms:** symptoms without any corresponding annotation data. These symptoms haven't been examined yet and need to be targeted next.

This information forms the basis for inferring a ranked list of likely diseases, but there are more factors with influence on the ranking: e.g. information about leading-symptoms, the intensity or novelty of symptoms, the age- and sex-specific incidence proportion⁷ of a disease and risk-factors due to other diseases or life-style (e.g. smoking). Most of this knowledge is also contained in Herold's "Innere Medizin" and we included such knowledge with the help of `owl:DatatypeProperties`.

The weighting of the different factors is still a topic of the ongoing discussions with our clinical partners, but the basic idea is to measure how well the patient's symptom information matches the typical symptomatology of some given disease in terms of precision and recall. Additionally, we provide the clinician with a "good graphical overview" about likely diseases with drill-down possibilities.

5.2 Scope and Features of the Prototype

In the first implementation we focused on five diseases, for which we already have annotated patient data from the MEDICO project: Hodgkin lymphoma, non-Hodgkin lymphoma, reactive lymphadenitis, colorectal carcinoma and diverticulitis. In interviews with clinicians we identified about 40 symptoms with relations to these diseases. Additionally, the clinicians listed so called leading symptoms (cardinal symptoms) for each disease.

As mentioned above the absolute and relative amount of present and absent symptoms is most important for creating a ranked list of likely diseases, where leading symptoms are weighted stronger. In this demonstrator we used a stacked bar diagram to show the absolute and relative amount of present and absent symptoms as well as leading symptoms (see figure 5). Information about risk-age and adapted incidence proportion is given additionally. Through hovering over the chart the clinician can see lists of the present, open and absent symptoms. Another tab gives a ranked list of open symptoms helping the clinician to plan next examinations.

5.3 Clinical Evaluation

In interviews with clinicians we got a positive feedback: inferring likely diseases shows the usefulness of annotations for clinical decision support. The graphical visualization and especially the possibility to get an overview of open symptoms helps to plan further examinations. Based on the prototype implementation the clinicians could explain what questions should be addressed next. Due to their feedback it would be of high relevance to include information about time-sequences of symptoms in order to detect their development, see differences between examinations and avoid inclusion of too old patient-data. Further, they see a need in representing the intensity of present symptoms and e.g. the amount of enlarged lymph nodes (1, 2, many). Additionally, the clinicians pointed out it would be extremely helpful to create a connection between medication and symptoms. Relying on the subclass hierarchy of medical ontologies is problematic: on the one hand there still exist simple errors like

⁷ The incidence proportion is the number of new cases within a specified time period divided by the size of the population initially at risk (Source: Wikipedia)

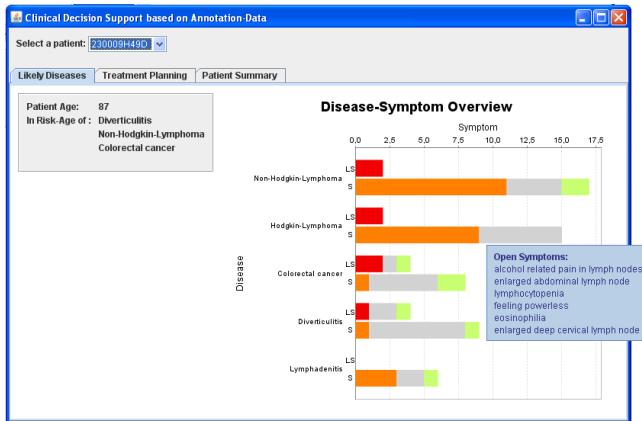


Fig. 5. Prototype implementation with graphical overview of present, open and absent symptoms and leading symptoms.

underlined in (Rector *et al.*, 2011), on the other hand – subclasses do not always represent their superclass in a *symptom-relevant* sense. E.g. in SNOMED CT "intended weight-loss" is a subclass of "weight-loss" (a symptom for lymphoma). Patient data annotated with the concept "intended weight-loss" will create a triple ":patient :showsSymptom :weight-loss". We cannot avoid that in a general way as we will not analyse all subclasses of referenced symptoms.

5.4 Revised Knowledge Requirements

We learned from the evaluation that we have to adjust our querying component in order to keep track of the time-sequence of the symptoms. This can be done easily as this information is contained in the MEDCIO-annotations, which are structured after studies with time-stamps. Basically we have to replace the relation "patient showsSymptom S" by a blank node construction holding also a date, the finding id, intensity etc. found in the annotations. Especially the finding id could be used to address the subclass-problem mentioned above, because this helps to track how we have come to specific present symptoms. In summary we set the following requirements for an enhanced Disease-Symptom-Ontology and application making use of it for clinical decision support:

- **Temporal information** is highly important in clinical diagnosis. The model has to represent the temporal relevance of symptoms, when and how long a symptom was present. This should allow to track the development of symptoms.
- **Inconsistency handling:** situations where annotations are contradicting each other (report: "no enlarged lymph nodes", image: "enlarged mediastinal lymph node") have to be presented to the clinician in an adequate way. Inconsistent sets of present and absent symptoms are hints for a deeper inspection.
- Represent a more fine-grade **significance** values of symptoms in addition to "leading symptoms" and the **intensity** of symptoms (1, 2, many enlarged lymph nodes).
- **Extend the coverage** of the model to examinations and medication.

6 CONCLUSION

Our approach of building a initial ontology referencing to the relevant medical ontologies is well suited to understand annotations with respect to symptoms and diseases. The description of symptoms with the help of links makes our ontology flexible for further applications. However through evaluation we identified several necessary enhancements, concerning the coverage and expressivity of the model. In future work we will address the listed requirements. Our long term aim is to build a generic context model for a more flexible and context-dependent interpretation of annotations under consideration of different medical background knowledge as illustrated in figure 1.

ACKNOWLEDGEMENTS

This research has been supported in part by the THESEUS Program in the MEDICO Project, which is funded by the German Federal Ministry of Economics and Technology under grant number 01MQ07016. The responsibility for this publication lies with the authors.

REFERENCES

- Channin, D. S., Mongkolwat, P., Kleper, V., Sepukar, K., and Rubin, D. L. (2010). The caBIG™ Annotation and Image Markup Project. *J. Digital Imaging*, **23**(2), 217–225.
- K, V. and Singaraju, J. (2011). Article: Decision Support System for Congenital Heart Disease Diagnosis based on Signs and Symptoms using Neural Networks. *International Journal of Computer Applications*, **19**(6), 6–12. Published by Foundation of Computer Science.
- Kulikowski, C. A. and Weiss, S. M. (1982). Representation of expert knowledge for consultation: The CASNET and EXPERT projects. In P. Szolovits, editor, *Artificial Intelligence in Medicine*, pages 21–55. Westview Press.
- Möller, M., Regel, S., and Sintek, M. (2009). RadSem: Semantic Annotation and Retrieval for Medical Images. In *Proceedings of the 6th Annual European Semantic Web Conference*. o.A.
- Opitz, J., Parsia, B., and Sattler, U. (2010). Evaluating Modelling Approaches for Medical Image Annotations. *CoRR*.
- Rector, A., Brandt, S., and Schneider, T. (2011). Getting the foot out of the pelvis: Modeling problems affecting use of SNOMED CT hierarchies in practical applications. *JAMIA*. **2011**;18: 432-440.
- Rector, A. L., Stevens, R., and Drummond, N. (2008). What Causes Pneumonia? The Case for a Standard Semantics for "may" in OWL. In *OWLED*.
- Rubin, D. L., Mongkolwat, P., Kleper, V., Supekar, K., and Channin, D. S. (2008). Medical Imaging on the Semantic Web : Annotation and Image Markup. *AAAI Spring Symposium Series Semantic Scientific Knowledge Integration*.
- Scheuermann, R. H., Ceusters, W., and Smith, B. (2009). Toward an ontological treatment of disease and diagnosis. *Summit on Translat Bioinforma*, **2009**, 116–120.
- Seifert, S., Barbu, A., Zhou, S. K., Liu, D., Feulner, J., Huber, M., Suehling, M., Cavallaro, A., and Comaniciu, D. (2009). Hierarchical parsing and semantic navigation of full body CT data. *Proceedings of SPIE*, pages 725902–725902–8.
- Seifert, S., Kelm, M., Möller, M., Mukherjee, S., Cavallaro, A., Huber, M., and Comaniciu, D. (2010). Semantic Annotation of Medical Images. In *Proceedings of SPIE Medical Imaging, February 13-18, San Diego, CA, United States*. o.A.
- Shortliffe, E. H. (1976). *Computer-Based Medical Consultations: MYCIN*, volume 85. American Elsevier.
- Wennerberg, P., Zillner, S., Möller, M., Buitelaar, P., and Sintek, M. (2008). KEMM: A Knowledge Engineering Methodology in the Medical Domain. In C. Eschenbach and M. Gruninger, editors, *5th International Conference on Formal Ontology in Information Systems. Formal Ontology in Information Systems (FOIS-08), October 31 - November 3, Saarbrücken, Germany*, volume 183 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, Amsterdam.
- Zillner, S. (2009). Towards the Ontology-based Classification of Lymphoma Patients using Semantic Image Annotations. In *Proceedings of SWAT4LS Semantic Web Applications and Tools for Life Sciences (SWAT4LS), Amsterdam, Netherland, November, 2009*.