AUTOMATIC SPATIAL PLAUSIBILITY CHECKS FOR MEDICAL OBJECT RECOGNITION RESULTS USING A SPATIO-ANATOMICAL ONTOLOGY

Manuel Möller, Patrick Ernst, Andreas Dengel
German Research Center for Artificial Intelligence (DFKI) and *University of Kaiserslautern, Germany*
{firstname.lastname}@dfki.de

Daniel Sonntag
German Research Center for Artificial Intelligence (DFKI), Saarbrücken, Germany
{firstname.lastname}@dfki.de

Keywords: medical imaging, semantic technologies, spatial reasoning, formal ontologies

Abstract: We present an approach using medical expert knowledge represented in formal ontologies to check the results of automatic medical object recognition algorithms for spatial plausibility. Our system is based on the comprehensive Foundation Model of Anatomy ontology which we extend with spatial relations between a number of anatomical entities. These relations are learned inductively from an annotated corpus of 3D volume data sets. The induction process is split into two parts. First, we generate a quantitative anatomical atlas using fuzzy sets to represent inherent imprecision. From this atlas we then abstract the information further onto a purely symbolic level to generate a generic qualitative model of the spatial relations in human anatomy. In our evaluation we describe how this model can be used to check the results of a state-of-the-art medical object recognition system for 3D CT volume data sets for spatial plausibility. Our results show that the combination of medical domain knowledge in formal ontologies and sub-symbolic object recognition yields improved overall recognition precision.

1 INTRODUCTION

During the last decades a great deal of effort went into the development of automatic object recognition techniques for medical images. Today a huge variety of available algorithms solve this task very well. The precision and sophistication of the different image parsing techniques have improved immensely to cope with the increasing complexity of medical imaging data. There are numerous advanced object recognition algorithms for the detection of particular objects on medical images. However, the results of the different algorithms are neither stored in a common format nor extensively integrated with patient and image metadata.

At the same time the biomedical informatics community managed to represent enormous parts of medical domain knowledge in formal ontologies. Today, comprehensive ontologies cover large parts of the available taxonomical as well as mereological (part-of) knowledge of human anatomy.

With the shift to the application of digital imaging techniques for medical diagnosis, such as CT, MRI, etc., the volume of digital images produced in modern clinics increased tremendously. Our clinical partner, the University Hospital Erlangen in Germany, has a total of about 50 TB of medical images. Currently they have about 150,000 medical examinations producing 13 TB of data per year.

To cope with this data increase (semi-)automatic image segmentation and understanding techniques from computer vision are applied to ease the task of radiological personnel during image assessment and annotation. However, these systems are usually based on statistical algorithms. Thus, detection and localization of anatomical structures can only be performed with limited precision or recall. The outcome is a certain number of incorrect results.

Our approach is to augment medical domain ontologies and allow for an automatic detection of anatomically implausible constellations in the results of a state-of-the-art system for automatic object recognition in 3D CT scans. The output of our system also provides feedback which anatomical entities are most likely to have been located incorrectly. The necessary spatio-anatomical knowledge is learned from a

large corpus of annotated medical image volume data sets. The spatial knowledge is condensed into a digital anatomical atlas using fuzzy sets to represent the inherent variability of human anatomy.

Our main contributions are (i) the inductive learning of a spatial atlas of human anatomy, (ii) its representation as an extension of an existing biomedical ontology, and (iii) an application of this knowledge in an automatic semantic image annotation framework to check the spatio-anatomical plausibility of the results of medical object recognition algorithms. Our approach fuses a statistical object recognition and reasoning based on a formal ontology into a generic system. In our evaluation we show that the combined system is able to rule out incorrect detector results with a precision of 85.6% and a recall of 65.5% and can help to improve the overall performance of the object recognition system.

2 RELATED WORK

Our primary source of medical domain knowledge is the Foundational Model of Anatomy (FMA) (Rosse and Mejino, 2007), the most comprehensive formal ontology of human anatomy available. However, the number of spatial relations in the FMA is very limited and covers only selected body systems (Möller et al., 2009). Thus, our approach is to infer additional spatial relations between the concepts defined in the FMA by learning from annotated medical volume data sets.

In (da Luz et al., 2006) the authors describe a hybrid approach which also uses metadata extracted from the medical image headers in combination with low-level image features. However, their aim is to speed up content-based image retrieval by restricting the search space by leveraging metadata information.

The approach in (Hudelot et al., 2008) is complementary to our work in so far as the authors also propose to add spatial relations to an existing anatomical ontology. Their use case is the automatic recognition of brain structures in 3D MRI scans. However, they generate the spatial relations manually, while a major aspect of our approach is the automatic learning from a large corpus.

Quantitative spatial models are the foundation of digital anatomical atlases. Fuzzy logic has been proven as an appropriate formalism which allows for quantitative representations of spatial models (Bloch, 2005). In (Krishnapuram et al., 1993) the authors expressed spatial features and relations of object regions using fuzzy logic. In (Bloch and Ralescu, 2003) and (Bloch, 1999b) the authors describe generalizations of this approach and compare different options to express relative positions and distances between 3D objects with fuzzy logic.

3 SYSTEM ARCHITECTURE

Figure 1 shows an abstraction of the distributed system architecture. It is roughly organized in the order of the data processing horizontally from left to right. All parsing results are stored in a custom-tailored spatial database.

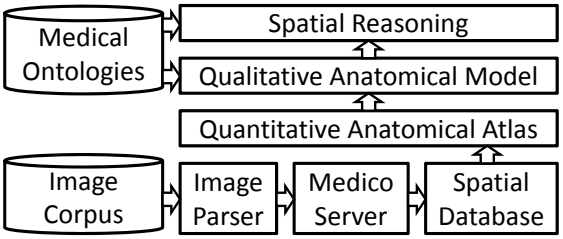

Figure 1: System architecture overview

3.1 Image Parser

To represent the results of the automatic object recognition algorithms in the format of our ontology we had to integrate rather disparate techniques into a hybrid system. The automatic object recognition performs an abstraction process from simple low-level features to concepts represented in the FMA.

For automatic object recognition we use a state-of-the-art anatomical landmark detection system described in (Seifert et al., 2010). It uses a network of 1D and 3D landmarks and is trained to quickly parse 3D CT volume data sets and estimate which organs and landmarks are present as well as their most probable locations and boundaries. Using this approach, the segmentation of seven organs and detection of 19 body landmarks can be obtained in about 20 seconds with state-of-the-art accuracy below 3 millimeters mean mesh error and has been validated on 80 CT full or partial body scans (Seifert et al., 2010).

The image parsing algorithm generates two fundamentally different output formats: *Point3D* for landmarks and *Mesh* for organs. Apart from their geometric features, they always point to a certain anatomical concept which is hard-wired to the model that the detection/segmentation algorithm has used to generate them. A landmark is a point in 3D without spatial extension. Usually it represents an extremal point of an anatomical entity with a spatial extension. Sometimes these extremal points are not part of the official FMA. In these cases we modeled the respective concepts as

described in (Möller et al., 2009). In total we were able to detect 22 different landmarks from the trunk of the human body. Examples are the bottom tip of the sternum, the tip of the coccyx, or the top point of the liver.

Organs, on the contrary, are approximated by polyhedral surfaces. Such a surface, called *mesh*, is a collection of vertices, edges, and faces defining the shape of the object in 3D. For the case of the urinary bladder, the organ segmentation algorithm uses the prototype of a mesh with 506 vertices which are then fitted to the organ surface of the current patient. Usually, vertices are used for more than one triangle. Here, these 506 vertices form 3,024 triangles. In contrast to the Point3D data, meshes are used to segment organs. For our test, the following organs were available: left/right kidney, left/right lung, bladder, and prostate.

3.2 Medico Server

In Figure 1 we show the overall architecture of our approach for integrating manual and automatic image annotation. One of the main challenges was to combine the C++ code for volume parsing with the Java-based libraries and applications for handling data in Semantic Web formats. We developed a distributed architecture with the *MedicoServer* acting as a middleware between the C++ and Java components using CORBA (Object Management Group, 2004).

3.3 Spatial Database

As we have seen in the section about the image parsing algorithms, the automatic object recognition algorithms generate several thousand points per volume data set. Storage and efficient retrieval of this data for further processing made a spatial database management system necessary. Our review of available open-source databases with support for spatial data types revealed that most of them now also have support for 3D coordinates. However, the built-in operations ignore the third dimension and thus yield incorrect results, e. g., for distance calculations between two points in 3D. Eventually we decided to implement a light-weight spatial database supporting the design rationales of simplicity and scalability for large numbers of spatial entities.

4 CORPUS

The volume data sets of our image corpus were selected primarily by the first use case of MEDICO which is support for *lymphoma* diagnosis. The selected data sets were picked randomly from a list of all available studies in the medical image repositories of the University Hospital in Erlangen, Germany. The selection process was performed by radiologists at the clinic. All images were available in the Digital Imaging and Communications in Medicine (DICOM) format, a world wide established format for storage and exchange of medical images (Mildenberger et al., 2002).

volume data available in total	777 GB
number of distinct patients	377
volumes (total)	6,611
volumes (modality CT)	5,180
volumes (parseable)	3,604
volumes (w/o duplicates)	2,924
landmarks	37,180
organs	7,031

Table 1: Summary of corpus features

Table 1 summarizes major quantitative features of the available corpus. Out of 6,611 volume data sets in total only 5,180 belonged to the modality CT which is the only one currently processible by our volume parser. Out of these, the number of volumes in which at least one anatomical entity was detected by the parser was 3,604. This results from the rationale of the parser which was in favor of precision and against recall. In our subsequent analysis we found that our corpus contained several DICOM volume data sets with identical Series ID. The most likely reason for this is that an error occurred during the data export from the clinical image archive to the directory structure we used to store the image corpus. To guarantee for consistent spatial entity locations, we decided to delete all detector results for duplicate identifiers. This further reduced the number of available volume data sets to 2,924.

Controlled Corpus

Due to the statistical nature of the object detection algorithm used for annotating the volume data sets, we have to assume that we have to deal with partially incorrect results. Hence, we decided to conduct manual corpus inspections using a 3D detect result visualization. The goal was to identify a reasonable set of controlled training examples suitable for generation and evaluation of a quantitative anatomical atlas and a qualitative model. These manual inspections turned out to be very time consuming. For each volume in the corpus a 3D visualization had to be generated and

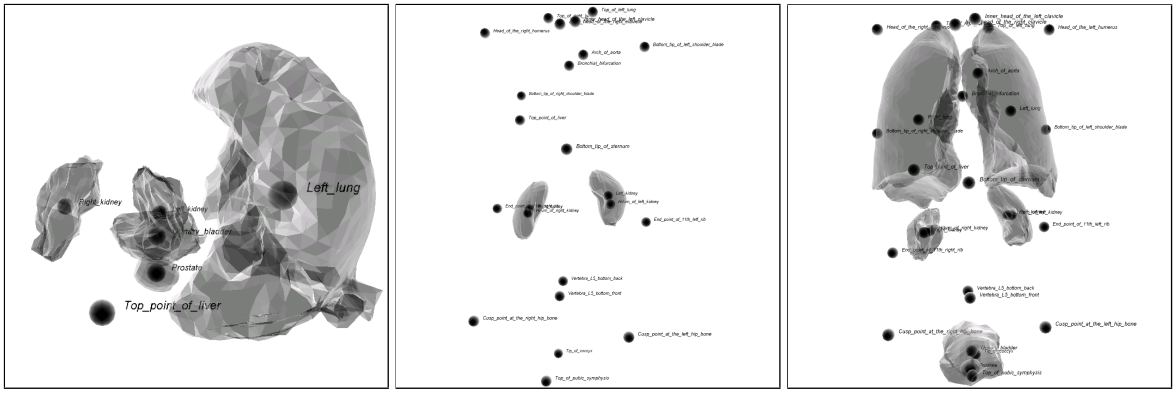

(a) (b) (c)

Figure 2: Visualizations of detector results: (a) incorrect; (b) sufficient; (c) perfect

manually navigated to verify the correct location of landmarks and organ meshes. After some training we were able to process approximately 100 volume data sets per hour. For higher accuracy, all manual inspection results were double checked by a second person resulting in a bisection of the per-head processing rate to about 50 per hour.

During our inspection we found that the quality of the detector results exhibits a high variability. Subsequently, we distinguish three quality classes: clearly incorrect, sufficiently correct, and perfectly correct. The visualizations in Figure 2 show one example for each class.

To have a solid basis for the generation of the spatio-anatomical model we decided to label a reasonable subset of the available volume data sets manually. We ended up with more than 1,000 manually labeled volume data sets. Table 2 summarizes the results quantitatively. All quantitative evaluations of the performance of the spatial consistency check are based on this corpus.

detector results inspected in total	1,119
apparently incorrect volume data sets	482 (43%)
sufficiently correct detector results	388 (34%)
perfect detector results	147 (13%)
volumes containing meshes	946 (85%)
volumes containing landmarks	662 (59%)

Table 2: Summary of the manual corpus inspection

We consider a detector result *incorrect* if a spatial entity configuration has been detected that is clearly contradictory to human anatomy. Figure 2 (a) shows such an example with arbitrarily deformed lungs. Normally, the lungs should be located vertically at about the same level. Here, this is not the case. Additionally, the prostate has been located on the top right side of the right lung although it belongs to a completely different body region.

A detector result is considered as *sufficiently correct* if it contains a reasonable number of landmarks and/or meshes. The following flaws distinguish them from perfect detector results (at least one condition is met): (i) It contains either only landmarks or only meshes. (ii) A minor number of anatomical entities has been detected at slightly incorrect positions. (iii) The overall number of detected anatomical entities in the detector result is rather low.

A *perfectly correct* detector result has to contain both landmarks and meshes. In addition, none of the landmarks or meshes is allowed to be located incorrectly. The anatomical atlas is learned only from detector results labeled as either sufficiently or perfectly correct. Incorrect detector results are discarded during model generation.

5 QUANTITATIVE ANATOMICAL ATLAS

Based on the spatial entities in the corpus we distinguish between two different types of relations to build up a quantitative atlas, namely: (i) *elementary relations* directly extracted from 3D data and represented as fuzzy sets, and (ii) *derived relations* which are defined using fuzzy logic and based on one or more elementary relations.

5.1 Elementary Relations

5.1.1 Orientation

The orientation or relative position of objects to each other is important to describe spatial coherencies. Typically, the fuzzy representation of the orientation depends on two angles used to rotate two objects on one another (Bloch, 1999a). The fuzzy set is thereby defined using six linguistic variables specifying the general relative positions: *above*, *below*, *left*, *right*, *in front of*, and *behind*. Their membership functions are basically the same.

$$\mu_{rel}(\alpha_1,\alpha_2) = \begin{cases} \cos^2(\alpha_1)\cos^2(\alpha_2) & \text{if } \alpha_{1,2} \in \left[-\frac{\pi}{2},\frac{\pi}{2}\right] \\ 0 & \text{otherwise} \end{cases}$$

They only vary in a direction angle denoting the reference direction, e. g., for *left* the angle is π.

$$\mu_{left}(\alpha_1,\alpha_2) = \mu_{rel}(\alpha_1 - \pi,\alpha_2)$$

More details about this approach can be found in (Bloch, 1999a). The definition of complex objects' relative positions is not straightforward. One possibility is to use centroids. Mirtich et al. describe a fast procedure for the computation of centroids (Mirtich, 2005). However, complex objects are reduced to single points and therefore information is lost. As the authors of (Berretti and Bimbo, 2006) state: "This still limits the capability to distinguish perceptually dissimilar configurations." For this reason we decided to use 3D angle histograms providing a richer quantitative representation for the orientation. A histogram H_A^R stores the relative number of all angles between a reference object R and a target A. The degree of membership is then obtained by computing the fuzzy compatibility between H_A^R and a particular directional relation μ_{rel}. Thus, we achieve a compatibility fuzzy set describing the membership degree.

$$\mu_{CP(\mu_{rel},H)}(u) = \begin{cases} 0 & \text{if } H_A^{R-1}(\alpha_1,\alpha_2) = 0 \\ \sup\limits_{(\alpha_1,\alpha_2),u=H_A^R(\alpha_1,\alpha_2)} & \mu_{rel}(\alpha_1,\alpha_2) \quad \perp \end{cases}$$

To compute a single value the center of gravity of $\mu_{CP(\mu_{rel},H)}$ is determined by

$$\mu_{rel}^R(A) = \frac{\int_0^1 u\mu_{CP(\mu_{rel},H)}(u)du}{\int_0^1 \mu_{CP(\mu_{rel},H)}(u)du}$$

Using this approach the orientation relations are now depending on the entire shape of the spatial entities. In addition, these histograms capture distance information. For example, if one object is moved closer or further away from another, the angles will also change according to the distance. Unfortunately, the membership degree computation is more complex compared to using centroids. However, since we are relying exclusively on the surface points of meshes, the computation time is acceptable with an average of 33 seconds for an entire volume.

5.1.2 Intersection

The detection of organ borders is a very difficult task in medical image understanding because it is mainly based on the tissue density (Bankman, 2000). However, adjacent organs can have very similar densities. Thus, detection is sometimes error-prone and objects may intersect. To check for such inconsistencies we are determining the degree of intersection between two spatial entities A and B. On that account, a new mesh or point is generated describing the intersection Int, so that the degree of intersection is determined by dividing the volume of Int with the minimum volume of A and B.

$$\mu_{int}(A,B) = \frac{V_{Int}}{\min\{V_A,V_B\}}$$

5.1.3 Inclusion

The inclusion of two spatial entities is similarly defined as the intersection. We say that a spatial entity B is included in an entity A, if

$$\mu_{inc}(A,B) = \frac{V_{Int}}{V_B}$$

Compared to intersection inclusion only considers the volume of the entity being included. For that reason this relation is not symmetrical contrary to all other relations described in this work.

5.2 Derived Relations

5.2.1 Adjacency

Many anatomical entities in the human body exist which share a common border or adjoin to each other, e. g., the border of the prostate and urinary bladder. These adjacent coherencies are represented using a trapezoid neighborhood measure depicted in Figure 3. Two spatial entities are fully neighbored if the distance between them is less than 2 millimeters. After that border the neighborhood decreases to a distance of 4.5 millimeters at which spatial entities are not considered as neighbored anymore. However, for an appropriate representation of adjacency the intersection between two objects has to be incorporated. This is important since if two spatial entities intersect, they

	$\top_{Goedel}$		$\top_{prod}$		$\top_{Lukas}$	
Relation	*avg*	*stddev*	*avg*	*stddev*	*avg*	*stddev*
Bronchial bifurcation, Right lung	0.0485	0.2001	0.0485	0.2001	0.0485	0.2001
Hilum of left kidney, Left kidney	0.0043	0.0594	0.0043	0.0594	0.0043	0.0594
Hilum of right kidney, Right kidney	0.0032	0.0444	0.0032	0.0444	0.0032	0.0444
Left kidney, Left lung	0.0427	0.1753	0.0427	0.1753	0.0427	0.1753
Left lung, Right lung	0.1556	0.3319	0.1556	0.3319	0.2617	0.3967
Left lung, Top of left lung	0.2322	0.3526	0.2322	0.3526	0.2322	0.3526
Prostate, Top of pubic symphysis	0.0116	0.0922	0.0116	0.0922	0.0116	0.0922
Prostate, Urinary bladder	0.2647	0.4035	0.2647	0.4035	0.7442	0.3408
Right kidney, Right lung	0.0376	0.1788	0.0376	0.1788	0.0383	0.1796
Right lung, Top of right lung	0.2900	0.3985	0.2900	0.3985	0.2900	0.3985
Right lung, Top point of liver	0.2288	0.3522	0.2288	0.3522	0.2288	0.3522
Top of pubic symphysis, Urinary bladder	0.0114	0.0918	0.0114	0.0918	0.0114	0.0918

Table 3: Mean values and standard deviations for fuzzy membership values for the adjacency relation depending on the choice of the t-norm

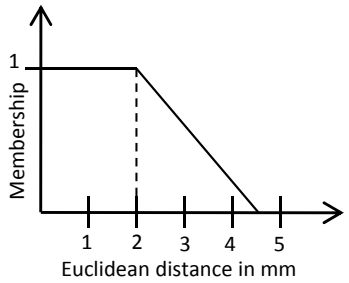

Figure 3: Graph of the fuzzy membership function for the linguistic variable *adjacent*

are not adjacent anymore. To formulate those circumstances using fuzzy sets, we comprised the degree of non-intersection and the neighborhood measure using a fuzzy t-norm (Klir and Yuan, 1994):

$$\mu_{adj}(A,B) = t[\mu_{\neg Int}(A,B), \sup_{x\in A}\sup_{y\in B} n_{xy}]$$

where the non-intersection is computed using the fuzzy logical *not*. Currently, three different t-norm based logic definitions are implemented, namely Lukasiewicz Logic, Gödel Logic, and Product Logic. The details of their definitions can be derived from (Klir and Yuan, 1994). Table 3 compares the average and standard deviations between the different logics. We decided to use the Lukasiewicz logic because it provides the highest average of actual adjacent concepts determined during a manual data examination. Additionally, the logic also yields the lowest standard deviations in comparison to the average value.

6 QUALITATIVE ANATOMICAL MODEL

Figure 4 illustrates our modeling of instantiated fuzzy spatial relations. It is loosely oriented on the formalism in the FMA for storing spatial relations. However, the value for each spatial relation is stored separately. Another difference is the representation with a term further qualifying the relation together with a truth value in a separate instance. Currently, we integrate orientation and adjacency in a qualitative model.

In order to create a qualitative anatomical model, we extracted instances containing the spatial relations described in Sect. 5. An instance describes the relation between two spatial entities occurring in a volume data set. To transform a relation into the model, a truth value is computed representing the mean of all extracted values of this relation. Thereby, the orientation is stated using a directional term, i. e., *left*, *right*, *in front*, etc. determined by the linguistic variables. On the other hand, the adjacency only gets a simple boolean qualifier. We determined a threshold of 0.2 (see Table 3) to distinguish between adjacent and not adjacent.

7 EVALUATION AND RESULTS

When an actual detector result is to be checked against the generic qualitative anatomical model, we first represent all its inherent spatial relations using the same formalism that we use for the generic anatomical model. This yields a set of OWL instances. Next, we iterate over all instances of the

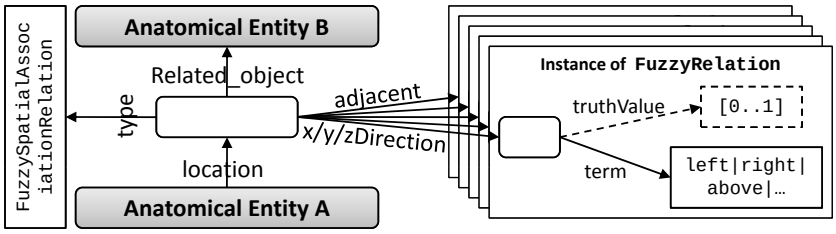

Figure 4: Illustration of the extended structure for storing the six linguistic variables separately and represent truth values in the ontology

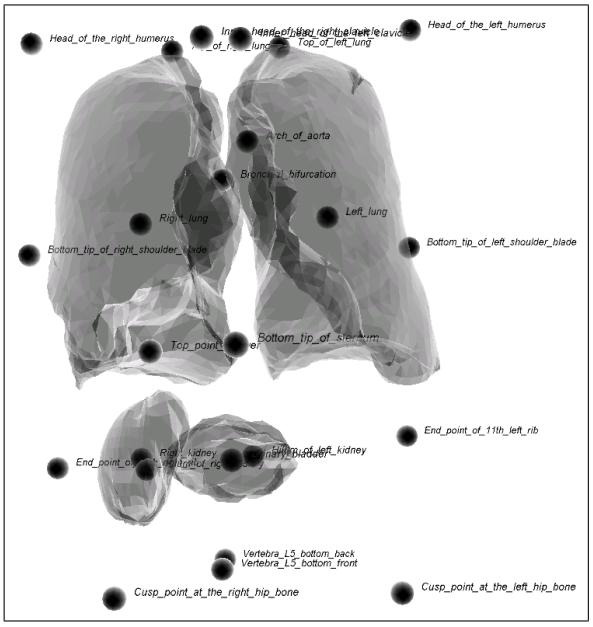

Figure 5: Visualization of the organ and landmark locations for an incorrect detector results (cf. the location of the urinary bladder)

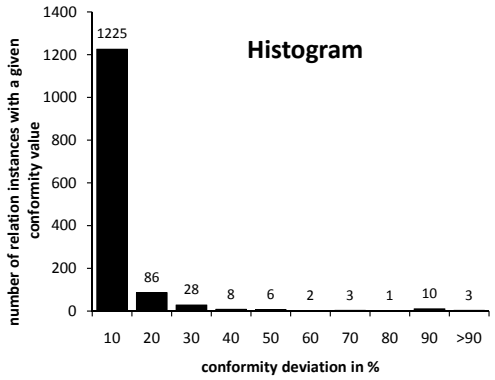

Figure 6: Distribution of the differences of the truth values between model and the detector result presented in Figure 5

detector result and compare their directions and truth values to the generic model. We consider a spatial relation instance to be *not conform with the model* if the truth values differ by at least 50%. We then count the occurrences of the anatomical concepts among the non-conform instances. The higher this number is for a given anatomical concept, the more likely the respective organ has been located incorrectly.

Figure 5 shows the visualization of an incorrect detector result. In the upper part you can see the two lungs and a number of landmarks. In the lower half you see one kidney and, to the right of the kidney, the urinary bladder has been located. This is clearly incorrect; in fact the urinary bladder should lie much further below. The other kidney has not been detected at all. Figure 6 shows a histogram of the differences in percent between the model and the spatial relation instances of this volume data set. Apparently, most of the relation instances have a comparably low difference to the model. Among all relation instances with a difference to the model of more than 50%, those with relation to urinary bladder account for 11 out of 16. This information gives evidence that the location of the urinary bladder is very likely to be incorrect.

Validation on Controlled Corpus

We performed a systematic evaluation of the spatial consistency check on our manually labeled corpus using four-fold cross evaluation. Our results show that the average difference in percent between the spatial relation instances in the learned model and the instances generated for an element from the evaluation set is an appropriate measure for the spatial consistency. The average difference to the truth value in the model for correct detector results was 2.77% whereas the average difference to the truth value in the model for incorrect detector results was 9%. Using 5% as a threshold to distinguish spatially consistent ($< 5\%$) from inconsistent ($>= 5\%$) yields a precision of 85.7% with a recall of 65.5% for the detection of spatially inconsistent detector results.

true positives	407
true negatives	431
false positives	67
false negatives	213
avg. difference correct detector results	2.7%
avg. difference incorrect detector results	9.0%
precision	85.7%
recall	65.5%

Table 4: Results of the spatial consistency check evaluation

8 CONCLUSION AND FUTURE WORK

We presented an approach fusing state-of-the-art object recognition algorithms for 3D medical volume data sets with technologies from the Semantic Web. In a two-stage process we augmented the FMA as the most comprehensive reference ontology for human anatomy with spatial relations. These relations were acquired inductively from a corpus of semantically annotated CT volume data sets. The first stage of this process abstracted relational information using a fuzzy set representation formalism. In the second stage we further abstracted from the fuzzy anatomical atlas to a symbolic level using an extension of the spatial relation model of the FMA.

In our evaluation we were able to show that this spatio-anatomical model can be applied successfully to check the results of automatic object detection algorithms. The detection of incorrect object recognition constellations can be performed with a high precision of 85.6% and a recall of 65.5%. The presented method can thus improve existing statistical object recognition algorithms by contributing a method to sort out incorrect results and increase the overall performance by reducing the number of incorrect results. Currently our anatomical model only covers directional information for pairs of spatial entities in our corpus. We plan to add spatial inclusion and intersection between entities.

Among our next steps is also a user evaluation of clinical applications making use of the reasoning, e. g., to support radiologists by suggesting anatomical concepts and relations during manual image annotation. Furthermore, our approach could be used to generate warnings for manually generated image annotations in case they do not conform to the spatial anatomical model. A clinical evaluation of these features is planned in the near future.

ACKNOWLEDGEMENTS

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

REFERENCES

Bankman, I., editor (2000). *Handbook of medical imaging: processing and analysis*. Elsevier.

Berretti, S. and Bimbo, A. D. (2006). Modeling spatial relationships between 3d objects. *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*, 1:119 – 122.

Bloch, I. (1999a). Fuzzy relative position between objects in image processing: new definition and properties based on a morphological approach. *International Journal Of Uncertainty, Fuzziness And Knowledge Based Systems*, 7:99–134. Notes.

Bloch, I. (1999b). On fuzzy distances and their use in image processing under imprecision. *Pattern Recognition*, 32(11):1873–1895.

Bloch, I. (2005). Fuzzy spatial relationships for image processing and interpretation: a review. *Image and Vision Computing*, 23(2):89–110.

Bloch, I. and Ralescu, A. (2003). Directional relative position between objects in image processing: a comparison between fuzzy approaches. *Pattern Recognition*, 36(7):1563–1582.

da Luz, A., Abdala, D. D., Wangenheim, A. V., and Comunello, E. (2006). Analyzing dicom and non-dicom features in content-based medical image retrieval: A multi-layer approach. In *Computer-Based Medical Systems, 2006. CBMS 2006. 19th IEEE International Symposium on*, pages 93–98.

Hudelot, C., Atif, J., and Bloch, I. (2008). Fuzzy spatial relation ontology for image interpretation. *Fuzzy Sets Syst.*, 159(15):1929–1951.

Klir, G. J. and Yuan, B. (1994). *Fuzzy sets and fuzzy logic: theory and applications*. Prentice-Hall, Inc. Upper Saddle River, NJ, USA.

Krishnapuram, R., Keller, J. M., and Ma, Y. (1993). Quantitative analysis of properties and spatial relations of fuzzy image regions. *IEEE Transactions on Fuzzy Systems*, 1(3):222–233.

Mildenberger, P., Eichelberg, M., and Martin, E. (2002). Introduction to the DICOM standard. *European Radiology*, 12(4):920–927.

Mirtich, B. (2005). Fast and accurate computation of polyhedral mass properties. *Graphics tools: The jgt editors' choice*.

Möller, M., Folz, C., Sintek, M., Seifert, S., and Wennerberg, P. (2009). Extending the foundational model of anatomy with automatically acquired spatial relations. In *Proc. of the International Conference on Biomedical Ontologies (ICBO)*.

Object Management Group, I. (2004). Common object request broker architecture: Core specification. online. Version 3.0.3.

Rosse, C. and Mejino, J. L. V. (2007). *Anatomy Ontologies for Bioinformatics: Principles and Practice*, volume 6, chapter The Foundational Model of Anatomy Ontology, pages 59–117. Springer.

Seifert, S., Kelm, M., Möller, M., Mukherjee, S., Cavallaro, A., Huber, M., and Comaniciu, D. (2010). Semantic annotation of medical images. In *Proc. of SPIE Medical Imaging*, San Diego, CA, USA.