Discussion of "Spatial-Symbolic Query Engine in Anatomy"

W.-T. Balke[1]; H. Handels[2]; I. Kalet[3]; M. Kimura[4]; C. A. Kulikowski[5]; L. A. Moura[6];
A. Pommert[7]; S. Schulz[8]

[1]Institute for Information Systems, Technische Universität Braunschweig, Braunschweig, Germany;
[2]Institute of Medical Informatics, Universität zu Lübeck, Lübeck, Germany;
[3]Radiation Oncology Department, University of Washington, Seattle, WA, USA;
[4]Medical Informatics Department, Hamamatsu University, School of Medicine, Hamamatsu, Japan;
[5]Department of Computer Science, Rutgers University, Piscataway, NJ, USA;
[6]Assis Moura eHealth, Porto Allegre, RS, Brazil;
[7]Department of Anatomy and Experimental Morphology, University Medical Center Hamburg-Eppendorf,
Hamburg, Germany;
[8]Institute of Medical Informatics, Statistics and Documentation, Medical University of Graz, Graz, Austria

With these comments on the paper "Spatial-Symbolic Query Engine in Anatomy", written by Antoine Puget and co-authors, with Dr. James Brinkley as senior author [1], *Methods of Information in Medicine* wants to stimulate a discussion on new forms of queries, considering anatomic knowledge. An international group of experts has been invited by the editor of Methods to comment on this paper. Each of the invited commentaries forms one section of this paper.

1. Comment by W.-T. Balke

The Semantic Web has matured into a Web of semantically meaningful applications. This is perhaps the most important insight that we can draw from the new paper "*Spatial-symbolic Query Engine in Anatomy*" presented by the Structural Informatics Group of University of Washington, Seattle [1]. Whereas the beginning of the Semantic Web era was marked by the search for a unified, consistent and thus in the logical sense truly semantic Web allowing for advanced reasoning tasks, practitioners in many domains have recognized the importance of cleverly putting together not necessarily consistent bits and pieces. In this way, easy to use and semantically meaningful solutions are ready to be delivered for important applications in a variety of domains.

This type of implementation – often referred to as mashup – tends to be technically simple, yet may be arbitrarily sophisticated in terms of usefulness. Here we see a SPARQL-based Web Service combining data from several sources to create a new application in anatomy as an important field of medicine. In particular, labeled 2-D axial images of the human body (the *Virtual Soldier* dataset) are combined with the *Foundational Model of Anatomy* (FMA) ontology to answer queries on spatial anatomical relationships. Indeed the applications for spatial anatomical queries like "Which vital organs are posterior, right-lateral or superior to the esophagus?" can be seen in many important areas ranging from emergency care to forensic medicine. The proposed method is a good example of easy and fast integration, where the major emphasis is not (yet) on efficiency, but on the creation or enrichment of valuable data sets that were not the original reason for producing the raw source data.

Indeed the enrichment of data sets or ontologies especially with spatial information is of increasing importance. This cannot only be seen by the current renaissance of geographic information systems (GIS) and location-based mashups especially for mobile applications, but also by the increasingly used notion of *knowledge spaces*. Knowledge spaces are abstractions that use spatial relationships to encode the real or perceived distance between real world entities or abstract concepts and have proved to be useful in several domains especially for visualization and personalization tasks. In this light more general applications like for instance connections of statistics in clinical trials or complex interdependencies between medical subjects in medical digital libraries may benefit from some of the techniques presented in the paper. Hence, the topic may even be of broader interest to the medical community than the anatomical application implies.

However, as already the preliminary evaluation with two domain experts clearly shows, there still is the important topic of *uncertainty* that has to be incorporated into the system in order to really fit the task. The authors' experimentation clearly shows that the two domain experts asked to provide a gold standard for the system's precision/recall analysis vastly differ in answering the same queries relying on their personal experience. In fact, the inter-rater agreement is rather limited, even more than shown by the actual evaluation of Cohen's Kappa that seems to be biased by the high number of true negatives relative to the other answer categories. Whereas one expert is rather consistent in both high precision and recall values, the other expert states mostly low precision values, but an almost perfect recall. This may give rise to speculations why two domain experts in anatomy can arrive at such vastly differing opinions.

Uncertainty may originate from three different sources: first t*he individual traits of the data set*, second the *semantic meaning of the query*, and third the *individual understanding of the user*. In the experiments im-

ages of a single human individual were used for deriving spatial relationships, whereas expert opinions will typically rely on textbook knowledge gained by observing spatial relationships over a diverse population of many individuals. Moreover, spatial relationships are usually not absolute, but may be partially true. Hence, query processing relied on the notion of "predominant" spatial relationships. Whether predominant is interpreted by human experts as the relative volume of a structure being involved in that relationship, the degree to which the relationship is satisfied, the relative importance of the structure regarding the relationship, or in some entirely different way remains doubtful. Finally, experts tend to form an individual view about a topic, often even with respect to the specific application. If for some application a certain relationship is often observed (e.g., in emergency care a certain organ maybe hurt quite often in accidents of a specific spatial type, even though there may be only a small spatial overlap), it may become a predominant spatial relationship for this user, however not for a typical user from for instance pathology.

Catering for these sources of uncertainty will be an interesting and valuable aim when researching spatial information systems in medicine in more detail. The degree of how uncertainty of either kind will influence a system's performance to a large degree is bound to depend on the semantic complexity of the system's task. Probabilistic databases featuring possible world semantics may be one solution, using knowledge based systems with uncertain reasoning methods may be an alternative, but not less promising route. In any case, the fresh mashup approach of the spatial-symbolic query engine raises hopes for a broad variety of useful applications in different areas of medicine.

2. Comment by H. Handels

2.1 Introduction

In the paper "Spatial-symbolic Query Engine in Anatomy" [1] the authors present an interesting approach to extract spatial-symbolic knowledge from medical images,

which is used in combination with the Foundational Model of Anatomy (FMA) ontology to get a symbolic representation of human anatomy. It is an inspiring idea to use medical image computing methods for the automatic definition of spatial relations in a given segmented image data set and to represent the results in standardized spatial-symbolic terms using ontologies. The paper describes first steps and results in this direction at the interface between symbolic representation of anatomy and medical image computing and shows the potential of this promising approach in future.

In my comment I would like to focus on the aspects of image based anatomical knowledge extraction and individualization of spatial-symbolic relations. In the field of medical image computing powerful approaches and tools have been developed to extract and use spatial knowledge about the position, shape and spatial relationship of image structures like organs, tissues, vessels, tumors etc. in the last decade [2, 3]. These methods are mainly applied to support image based medical diagnostics and therapy. However, they also can play an important role for the generation of comprehensive and individualized formal descriptions of spatial anatomical relations in future. Therefore, I first give a brief description of the main developments in this field in the next chapter.

2.2 Extraction of Spatial Knowledge from Medical Images

Three-dimensional image data sets acquired in computer tomography or magnetic resonance imaging implicitly contain anatomical information about the positions of image structures like organs, tissues, vessels, tumors etc. and their spatial relationships. In the field of medical image computing, on the one hand 3D models of image structures are extracted from the image data and visualized in a pseudo-realistic way. These 3D visualizations can be used to improve diagnostics [4, 5] and therapy [6, 7] as well as medical education [8, 9]. On the other hand, in modern image analysis systems model based methods are often used during image segmentation and analysis to take anatomical knowledge

about the position and shapes of organs into account. Here, atlases that show the anatomy of a single individual as a reference segmentation [10, 11] as well as statistical shape models describing the organs' mean shapes and their shape variations in a group of individuals [12–14] have been developed to represent complex anatomical relationships and inter-individual variations in the human body. A main goal of these approaches is to improve the robustness and accuracy of the segmentation and recognition of image structures of interest by using anatomical knowledge, e.g. by atlas based segmentation and landmark propagation.

2.3 Generation of Spatial-symbolic Descriptions of Anatomy

In their paper, Puget et al. present a methodology to extract spatial relations of organs describing the relative position of image structures to each other from a 3D image data set. The goal is to represent this knowledge explicitly in terms of symbolic descriptions like posterior, anterior, left-lateral etc. For the automatic generation of this spatial-symbolic description the authors present a 2D-oriented method applying 2D projections of the 3D image objects to the considered 2D planes in a set of slice images. However, the nature of the spatial relationship between organs and other image structures is three-dimensional. Hence, it would be an interesting question whether the application of 3D image computing techniques could achieve symbolic-spatial descriptions with higher accuracy. Because these computations can be done in a pre-processing step, it should be possible to also use advanced 3D approaches to extract spatial-symbolic relations, automatically.

In comparison to the spatial information contained in a 3D image data set the symbolic description generated by the approach of Puget et al. is rough and reflects only a small part of the spatial information available. Advanced image processing methods enable to extract further information on spatial relations and quantitative spatial parameters, automatically [2, 3]. Therefore, the planned extension to provide distances between organs and their

centers of mass seems to be a first step in this interesting direction and will give the user more helpful qualitative and quantitative information.

2.4 3D Modeling and Visualization versus Symbolic Representation of Anatomy

I would like to support the authors' statement that the integration of 3D visualizations of anatomical structures would help the user significantly instead of just providing textual lists of spatial relations. However, the generation of anatomical 3D models that can be explored interactively (e.g. by translation, rotation, zooming of organs and the 3D scene etc.) is also a very intuitive alternative to give the user information about the position and spatial relationships of organs etc. and tissues. Hence, it would be interesting to see whether the implicitly modeled knowledge on spatial relationships coded in a virtual 3D scene of a patient body or atlas [8] is more helpful as the textual lists and explicit descriptions of the relationships that are generated by the presented system of the authors.

2.5 Towards the Representation of Patient-individual Anatomy

At the current state, the expected practical value of the presented system seems to be limited, because the representation of anatomical information is based on only one atlas data set, the so-called Virtual Soldier data set, which consists of label images with 437 segmented tissue structures and organs of the Visible Human data set [15]. The evaluation of the authors gives no impression of the practical value of the system. Only the accuracy of the generated spatial-symbolic descriptions extracted from one data set is considered and aspects like user acceptance and benefits for the user are not addressed.

However, obviously the value of the presented system would be increased strongly, if not only one reference data set, but the individual image data set of a patient can be used. This is motivated by the variability of anatomy and the occurrence of pathologies like tumors that can lead to individual changes in the spatial relationships of organs and tissue structures.

But the transfer of the methods described to individual patient data sets is challenging, because a pre-requisite of the application of the methods presented is that all image structures of interest are segmented and can be addressed by their names. However, the time needed for manually segmentation of all interesting image structures is not acceptable for the user and automatic, highly accurate and robust segmentation methods are currently not available. But advanced registration and segmentation techniques could be applied to generate a rough segmentation of image structures automatically, and it would be interesting to see whether the obtained results will be sufficient for the generation of different spatial-symbolic descriptions.

2.6 Conclusion

In the paper, Puget et al. addressed an interesting scientific field at the interface between symbolic representation of anatomy, ontologies and medical image computing. The authors describe first steps in this promising direction, but further developments of the system presented are desirable. Particularly, the integration of advanced image computing and visualization methods could lead to a significant improvement and extension of the possibilities to generate comprehensive spatial-symbolic descriptions of the human anatomy as discussed. In clinical practice, patient-specific relations and descriptions of anatomy based on an individual 3D image data set would be very helpful, but the complex structure of human anatomy on the one hand and its variability depending on age, sex, diseases etc. on the other hand will keep the automatic generation of patient-specific segmentations and spatial-symbolic descriptions with high accuracy and robustness being a challenge.

3. Comment by I. Kalet

The understanding of spatial relationships in the human body has always been important in medical practice, particularly in the practice of diagnosis, surgery and radiation therapy for cancer. As Puget et al. [1] point out, this knowledge is gained through the study of anatomy in medical school and by experience and training in the various specialties. However, the increased precision of surgical procedures and radiation treatment machinery has made this ever more challenging. The wealth of detail in digital medical images has opened up opportunities to combine quantitative and qualitative knowledge and methods to improve therapy. Thus, this work on formalizing and automating queries about spatial relationships and utilizing extensive knowledge resources such as the FMA is very important.

The authors envision two kinds of facilities. One is a query formalism and processing methods for expressing and automatically answering queries regarding spatial relationships among internal body parts (initially in this work, just at the granularity of organs) and the other is a well designed web application that uses the first facility to provide an easy user interface. My comments will be directed only to the first of these two visions, which is the more radical. The user interface focus is rather more conservative, preserving the idea that people look up answers to well defined questions and they then do the hard part themselves. The possibility of such a query engine as support for automated treatment design assistance is much more visionary and enticing.

The authors mention two applications of spatial reasoning to radiation therapy planning. One is to answer the question, "which lymph nodes are likely to have metastatic disease present?". This question is more about the topological properties of the lymphatic system, its connectivity, than about relative location in space [16]. Their second question, "which vital organs near a tumor can be affected by a planned radiation treatment?" is indeed about precisely the kind of relations described in their work. Such information can be used to drive an automated radiation therapy plan construction system that is based on rules, as described in very early work in this area [17, 18]. In Paluszynski's work, the spatial relationships were highly customized to the
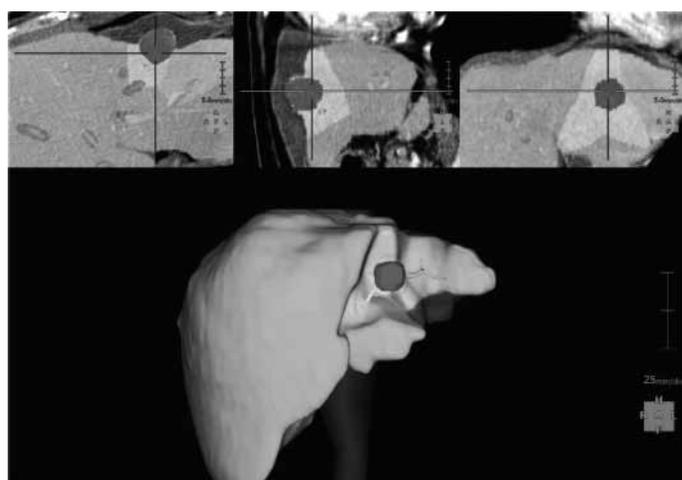
Fig. 1
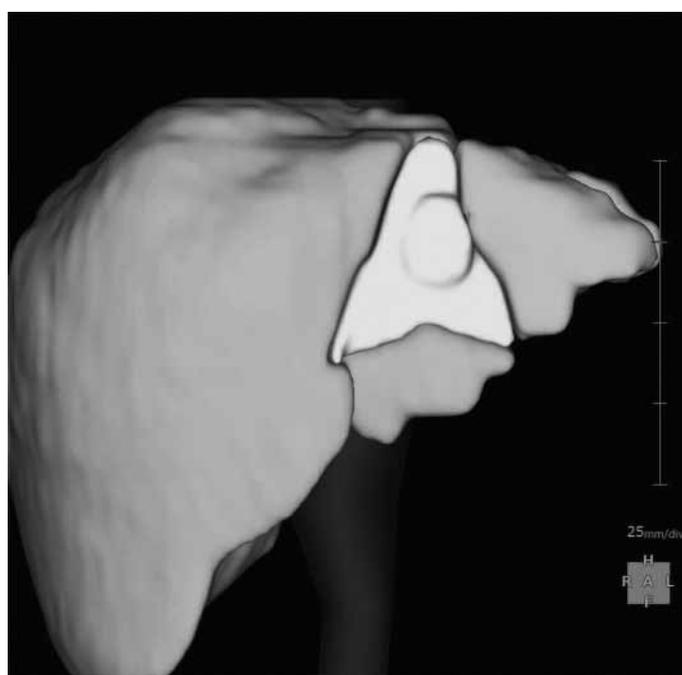Pointing the target tumor on three direction slices

Fig. 2
Lobes in which the tumor lies are illustrated

planning environment and did not use a reference anatomy model such as the FMA. Incorporating the FMA as proposed in the present work could be a major step forward on a very difficult problem.

Puget et al. have demonstrated the feasibility of automated spatial semi-quantitative reasoning about structures, and the incorporation of supporting knowledge from a symbolic knowledge resource such as the FMA. However, some challenges remain. They start with annotated images. This step is by itself a largely manual process. The state of the art of image segmentation for radiation therapy is still far from complete automation or even significant semi-automation for vital organs. For delineation of the target volume (the region including the tumor mass and immediate tissue likely to contain microscopic disease), there are no automated methods whatsoever. In fact there is not a "gold standard" or much agreement among clinicians about what the target volume should be for any given patient [19]. The focus for target volumes has been almost entirely on new ideas for generating images that show such volumes as segmentable entities with well de-

fined boundaries. This is an important matter, since the query about what vital organs are in the radiation beam path is really about what organs have the specified spatial relation with the target volume, not with any entity that exists in the FMA.

Scaling up is another open problem identified by the authors. Searching and analysing every relevant structure for a given query does not look feasible, but neither does the possibility of precomputing and caching the results of parts of the analysis. For the radiation therapy application, millimeter spatial precision is now critical in order to fully use the capabilities of the most advanced machinery. It is no longer possible to restrict the organs or substructures of concern to only a few, as was typically done in the past.

None of this should deter this work or further pursuit of the vision. It is only a perspective, to temper the wish for immediate payoff in terms of clinical application. The vision is a long term one, and this is a very important first step. I look forward to the next ones.

4. Comment by M. Kimura

Image processing cannot be accomplished only from one direction, i.e. bottom up by image data processing, and top down from hierarchical knowledge. Bottom up methods are based on classical thresholding, edge detection, contour definition and mass definition. These methods recognize each mass (organ, in medicine), but semantic notation can only be done by other kinds of knowledge, even after overcoming individual case variations. Top down methods, using hierarchical knowledge, tell us relations such as "part-of", but location knowledge is hard to be combined.

Tying this bi-directional approach with human expertise, as a hybrid system, has already been advanced to commercial products. An example of a liver cancer operation simulation is shown in ▶Figures 1–3. Based on thin slice CT images of the abdomen, the liver is recognized by each lobe. First, a physician as user points to the targeted tumor (▶Fig. 1) on three directions. Then lobes where the tumor lies are marked (▶Fig 2). If these lobes were re-

moved, a simulated image of the post-operation liver will be given. On the other hand, the artery and portal veins, which flow to and from the cancer lying lobes, are illustrated, in order to give the user surgeon the plans of which vessel to ligate (figures courtesy of Fuji Film Inc.).

This article [1] is new and interesting in that it introduces new dimension, spatial knowledge representation, additional to FMA ontology, as hierarchical knowledge. As far as this reports, results gained are enough, as this research is still open-ended. There are two directions after this, prepared answers to every possible spatial relation question, or each case based custom recognition, may be based on real individual CT or MRI images and for operation simulations. In any case, web-based open ended usage gives us a good starting point for future research.

For future research, two directions can be added to that which is described in the article. One is introducing connective relation information. Macro anatomy or radiology tell us skeletal structure, arm and legs, have less variations, then the head, chest. The most locational variations are shown by abdominal cavity organs. This will be because of skeletal cavity strength, and connection tightness-looseness. This information will add more preciseness to the rate of right answers to the spatial relation questions.

Another direction, mainly for real clinical use, is incorporating human aid, like in the above mentioned liver tumor application, to give milestone for the inference. The reason why the above mentioned liver tumor simulator needs pointing tumor by human is that error in this detection caused a fatal result, when used for real operation, even if the machine recognition results in 99%.

I encourage the authors, as well as readers, to further this research on these promising directions, to achieve real clinical use.

5. Comment by C. Kulikowski

The issues surrounding the design of effective and efficient query engines over both spatial and symbolic knowledge sources are fundamental from perceptual and cognitive perspectives, and are essential to help design human-machine systems to interpret the complex medical image datasets so essential for clinical practice.

In the paper by Puget et al. [1] from the University of Washington's pioneering research group on structural informatics, a prototype system is described for combining spatial and symbolic representations to answer queries about the localization of organs and their likely involvement after injury or disease. The work is an important step forward in specifying the kinds of informatics challenges that need to be overcome in order to use annotated 2-D cross-sectional images, so as to answer spatial queries about relative locations of target organs in relation to other organs that might involve pathology. The system goes beyond the traditional spatial-symbolic querying for image retrieval by generic similarities among images [20], by deriving its semantics from the medically-specific Foundational Model of Anatomy (FMA) [21]. This enables the system to match terms from the annotations of the organs in the 2D cross-sectional images, and help interpret relative spatial positions such as "anterior-to", "posterior-to", "superior-to", and "inferior-to". A specified target organ

of a query must be correlated to the query structure as seen in the images, and serves to divide up the surrounding space into the four major relative positions of anterior, posterior, above, and below, as well as sidedness positions like "antero-left-lateral". But, processing speed is a problem, even when restricted to a limited subset of the FMA, and a single dataset from the Virtual Soldier [22], with 437 segmented and labeled (annotated) anatomical structures. This is because there are a very large number of potential relative positions of intersections of projected anatomical objects that could yield positive answers to the relative localization queries. Despite an efficient cache implementation the number of potential projected intersections remains very large even when a significant amount of overlap (40%) is specified so as to eliminate many of the possibilities. As a consequence, the current implementation does not allow for real-time image processing, but does provide a proof of concept about how this kind of system can work.

The system is implemented using a web service model which can call on alternative sources of information for comparison beyond the FMA, as well as multiple experts, so as to obtain a consensus opinion on the definitions of spatial relationships between the anatomical objects. Inter-

Fig. 3
Artery and portal veins to and from the tumor lobes are illustrated.

action is through a graphical user interface (GUI) with capabilities for assessing the symbolic interpretation by visual mapping onto the image dataset, thus permitting effective mixed-mode evaluation. Integration with other semantic web applications will be enabled in this way. Since the system is still a prototype, the authors make clear that it is mainly a "framework for building end-user applications, rather than an end-user application in itself". However, it gives a clear illustration of just how detailed the specifications of relative spatial positioning need to be made in order to obtain answers to even simple comparative anatomical queries.

Results from a preliminary evaluation on a random sample of 10% of possible query structures are presented showing that good performance can be obtained in terms of low false positive and false negative rates for transverse and sagittal direct relations among anatomical objects. Since there is no "gold standard" available as a baseline for comparison, two experts carried out the evaluation, but this did show that inter-observer variability can be a significant factor, with one expert systematically avoiding false negatives for all spatial relationship types (keeping the rate mostly at zero or with at most one false negative case involving the inferior relationship) while tolerating a higher false positive rate, in contrast to the other expert, who balanced the precision and recall rates quite consistently. This inter-observer difference illustrates the complexity of designing comparative visual evaluation studies involving so many possible ways of defining spatial relationships between anatomical query objects and the possible targets in the image datasets. Greater standardization will be needed to scale up studies of this kind, though the variability between experts is only the tip of the iceberg in what is a long-standing problem of comparing features within and between objects in an image or scene where issues of perceptual grouping are important, and as yet far from being well understood for realistic visual stimuli [23].

Despite the difficulties reported, the authors provide a number of sensible suggestions for additional features that could improve the utility of their application system, making this preliminary study and evaluation a valuable reference for others seeking to develop this important type of medical spatial-symbolic system, so critical for interpreting and exchanging clinical image data across platforms and applications over the web.

6. Comment by L. A. Moura

Search engines have become so popular we have the impression they have always been around and the feeling that it is impossible to live without them. The kinds of search available cover from symbolic data, structured or not, to graphics and image content [24, 25] to samples of recorded music [26]. Databases worldwide store knowledge that allows humans and machines to search for music from signal recordings, pictures from samples, and places from pictures with or without GPS-referenced data. The amount of technology aboard our smartphones is huge and tailored to make the most of available services, yet at the same time services are designed to take advantage of available devices. Fast processors, GPS decoders, position sensors, compasses, 3- or 4-G communication, computer graphics, maps, charts, touchscreen displays, Wi-Fi, Bluetooth, HD video and photo cameras, voice synthesizers and language processing software are among the features most users expect any currently available model will offer. Needless to say that a decade or so ago most of this was "conceivable" but not "doable", and that the reality has, in many regards, beaten fiction.

Apart from being personal leisure and communication gadgets, mobile devices and the services they access have changed our lives and our work. Many activities as civil aviation, banking and shopping have been almost completely reinvented around the existing resources.

Although all measures of success applied to Health Care in the World Wide Web lead to very impressive figures (number of available services, specialized search engines, hits, unique users, product sales, client and investor's interest and so on), eHealth activities are still far behind the areas mentioned before, despite governments and private sector efforts and investment poured into eHealth around the world.

eHealth can and will change Health for better, but this is not a trivial task. The main reason for my having enjoyed so much reading **Spatial-symbolic Query Engine in Anatomy** [1] is that it tackles, in a very elegant way, some of the most vital concepts that are required to make eHealth mainstream, useful and sustainable.

To start with, the project stems from the need to provide answers to queries that are very simple from the human perspective, but complex from a computational viewpoint. The notions of "lateral", "anterior" and "inferior" are thus expressed in a very logical and coherent machine-compatible way. This is very much contemporary in that lay people and experts alike are willing to pose questions such as "a bullet in the chest is likely to hit which organs?" in a straightforward manner. In other words, the project hides the computational complexity behind what can be called a "simplicity layer": the interface that gives enhanced usability to the system.

The authors have built this project taking into account existing knowledge and, moreover, reusing existing standards and standardized models, methods and languages, as is the case of RDF, SparQL and OWL all of them W3C recommended standards [27]. Also the vocabularies and ontologies they use as the foundations for the work are published standards. All that means other groups and standard-based applications can collaborate to bring in new functionalities in an orchestrated way.

Finally, and a very important factor in my view, the authors worried themselves to offer their Query Engine as a web-service, which means that not only people, but machines throughout the world, can query the service and take advantage of it.

By being standards-based, using a service-oriented architecture and making complex stuff seem simple, this project forms an emblematic example of how the eHealth application of the future will be: simple to use, focused on a single specific and often complex subject, offered as a service – to humans and machines alike, ubiquitous, scalable and interoperable. Such are the features that will make applications part of an eHealth platform that seamlessly

integrates applications, knowledge and services, thus being able to be the engine that will help "run" the health system.

Several national and regional initiatives are on their way to build eHealth platforms that offer a common and flexible ground for applications to interoperate [28–30]. Of course, some of the founding building blocks are preferably offered at national or provincial levels. These include registries that provide unique IDs for patients, health care providers and health care organizations, for example [31].

In these days of mHealth, a plethora of mobile applications typically aiming at supporting vertical or local needs have been described that are not based on standards or even basic common building blocks as ICD-10, for instance. This is very worrying because such applications – if successful – will fragment information flow and, thus, fragment delivery of care [32].

On the other hand, it is not difficult to imagine series of "stand-alone" tasks being offered in the cloud, as services, in an interoperable fashion. Knowledge on drug-to-drug or food-to-drug interaction as well as guidelines and evidences can be offered as web-services instead of full-fledged applications, simplifying the business transactions as well as the management of such complex systems. The **Spatial-symbolic Query Engine in Anatomy** project we discuss, here, is a clear example of a very complex environment that requires specialized people to keep it up, running and evolving. By being offered as a service, more people can benefit from it, without having to cope with the inherent complexity.

The paper under appraisal is an example to be followed: it is innovative, focused on a single problem, built on existing foundations that are expanded by it; compliant with standards and best practices, offered as a service, and designed for interoperability with other systems. Indeed, an important contribution to our field.

7. Comment by A. Pommert

Is Canada north or south of the United States? Apart from Alaska, probably most of us would not hesitate to place it north of the US. However, if you happen to be in Detroit, things appear somewhat different, as the closest way to Canada is to go straight down south. So a more comprehensive answer to this question could be "mostly north".

Not surprisingly, things are even more complicated in the medical domain. The three-dimensional human anatomy is the most complex structure we know; its organs, organ parts and tissues are highly interwoven and contain myriads of descriptive problems like the one mentioned above. And yet, we are able to make statements about the relative positions such as "the heart is anterior to the lung". Or are we?

Anatomical terms of location such as anterior, (left) lateral or superior serve a very important purpose in medicine, as they provide an abstract (and thus somewhat simplified) description of spatial relations, which can easily be communicated. However, these terms are not well-defined in a mathematical sense. Is the descending thoracic aorta anterior to the lower lobe of the left lung? The answer to this question may also depend on the context, e.g. whether it is discussed in an anatomy class, or the intent is to direct a biopsy needle. In the latter case, even a very small portion in the way should result in a positive answer.

Descriptions of the human anatomy on a symbolic level, i.e. in terms of anatomical entities and their relations, have made a remarkable progress in recent years. The most comprehensive work, the Foundational Model of Anatomy (FMA), contains tens of thousands of entities and relations such as "part of" or "branch of". However, no comprehensive symbolic description of the spatial relations of these entities is available so far.

The paper from the Structural Informatics Group at the University of Washington [1] describes an important work to close this gap. Interestingly, the authors did not bother to define the spatial relations between the various anatomical entities manually, but automatically calculated them from a segmented three-dimensional dataset. For terms such anterior and posterior, rather strict definitions were used, based on projections of the structures parallel to the anterior-posterior axis (likewise for left lateral/right lateral and superior/inferior). This way, it could be determined how many percent of a structure are located e.g. anterior to another structure.

To answer questions about spatial relations, the authors developed a query engine which takes a structure, a spatial relation type and a threshold value x. The threshold limits the results to those structures which are involved with at least x percent. This mechanism allows to somewhat balance the answer on the context.

It is one of the merits of this work that the authors did not leave it at this. They also tested the results of the spatial-symbolic query engine against the opinions of two anatomists. One of the experts turned out to be consistently more critical with the resulting lists than the other. This may be due to the general problems of abstract spatial descriptions discussed above, the somewhat unsharp definition of terms such as anterior, or different assumptions about the context. It is an interesting question whether the percent threshold will be sufficient to handle these problems.

In many ways, this pioneering work can be considered as an important first step. Possible extensions include the coverage of natural variability, which in a first approximation requires little more than scanning other segmented datasets, or the inclusion of other spatial relations describing e.g. distances. For a full integration in the FMA, the combined semantics of spatial relations and patrimonies need to be understood.

The spatial-symbolic description of human anatomy clearly has the potential to become a powerful tool for many applications ranging from education to clinical systems informing about the possible consequences of injuries, diseases, or invasive procedures. Furthermore, the knowledge represented here may be used for an automatic segmentation of radiological images, or even an automatic detection of anomalies. Very likely, the obtained results will be useful for geographic information systems and in other non-medical domains with similar types of problems as well.

8. Comment by S. Schulz

In biomedical knowledge representation there is an increasing consensus about the need of ontological foundations, which
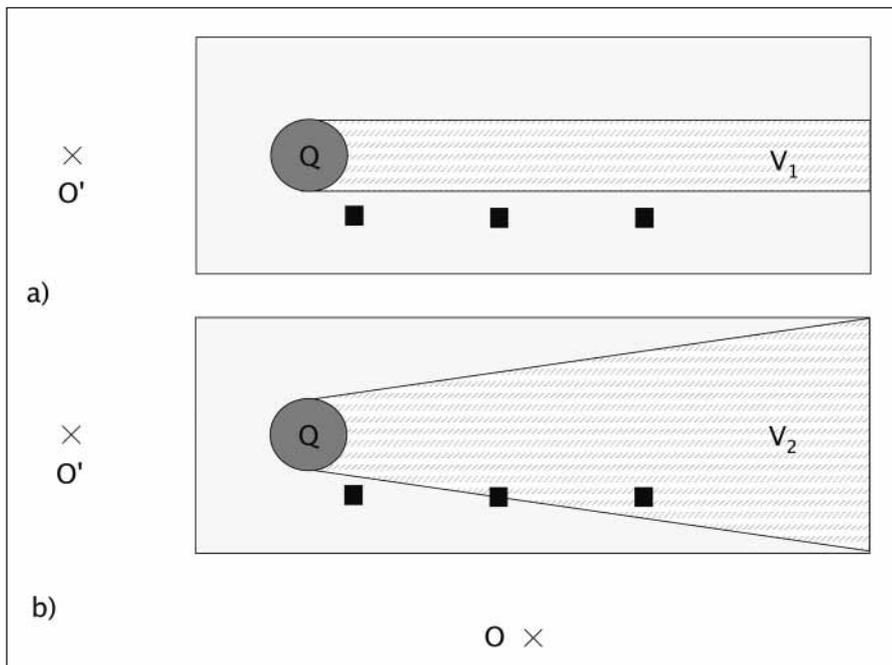
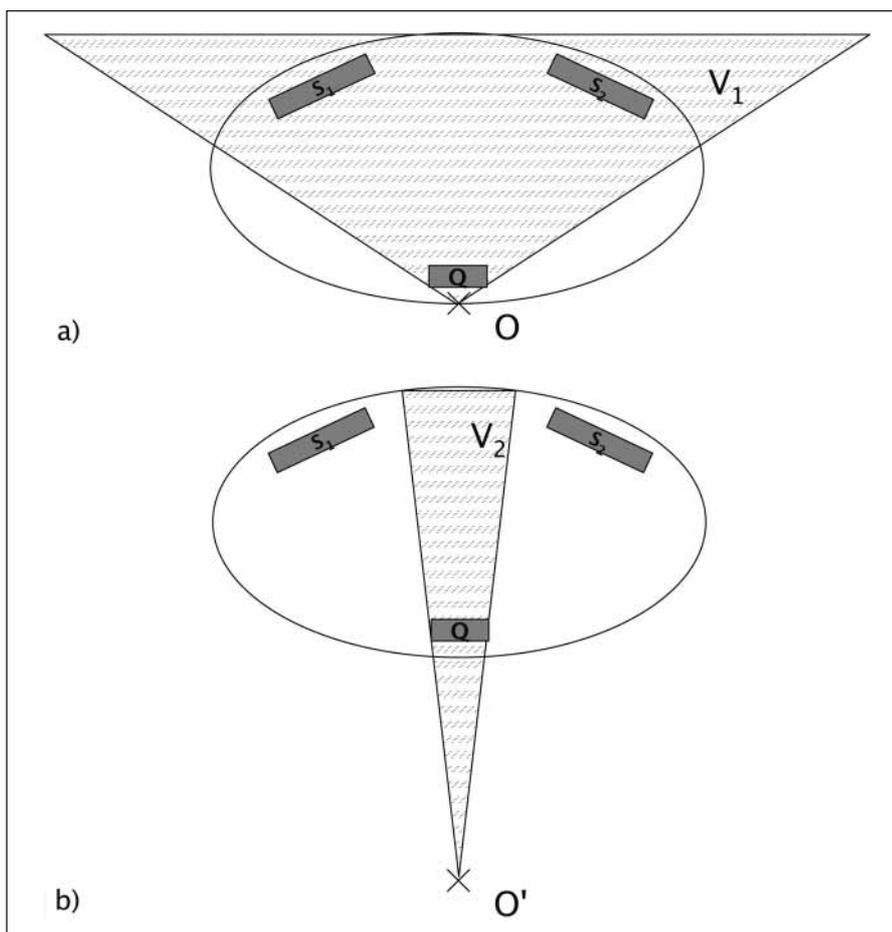**Fig. 4**   Parallel vs. conic projection volumes

Fig. 5   Conic projection volumes. The importance of the viewpoint

provide axiomatic statements of what can be taken as universally true for all entities that instantiate a given natural kind [33]. E.g. all cells have a membrane, all lungs have alveoli, and all arteries are blood vessels and have a wall and a lumen. Basic relations have been proposed to describe how entities are related with each other, independent from any observational context [34]. Examples of such relations are "has-part", "participant of", or "located in". Also the FMA uses such ontological relations, mostly to relate parts and wholes.

Locative relations such as expressed by "anterior", "posterior", "left", "right", or prepositions like "behind", "above", are not, a priori, ontological relations: Whether the heart appears in front of the esophagus or right to it depends on the position of the observer. This effect is eliminated if the spatial relations are normalized according to the standard anatomical position. As Puget et al. [1] point out, it would then suffice to assert the "classical" anatomical spatial relations like "anterior", "posterior" from volume data, according to strictly geometrical criteria. Does this represent the meaning of these terms in a cognitively adequate way?

Let us assume a scenario with three bullet fragments A, B, C with a diameter of 5 mm found in the thorax to the left of the descending aorta (Q), as depicted by ▶ Figure 4a. According to the standard anatomical position, "to the left" corresponds to "to the right" on the image. Following the criterion used for constructing query volumes by parallel projection as proposed by Puget et al., none of the three target structures would be inside the query volume for "to the left" and would instead be classified as "antero-left-lateral". However, we would not be surprised if most observers agreed with this only for the object A but not for C, which they would rather see as positioned left to Q.

Are there other ways to construct query volumes, which would do better justice to these distinctions? ▶ Figure 4b depicts an alternative solution. Here, only the target structure C would be fully inside the query volume for "to the left". This approach emphasizes the importance of the visual angle between two objects. Assuming O′ as the observer's position, the visual angle is

much greater for A with relation to Q than for B, as on the retina the distance between the images of Q and A is greater than the distance between Q and B. The position of A is therefore viewed more deviant from the position of Q. Equally, B and C would be perceived more closely related to an imaginary line between running from O′ through Q. Even more, as depicted by ▶Figure 4b, B would be partly and C fully eclipsed by Q when observed from O′.

The construction of conic instead of parallel query volumes could therefore be a more cognitively adequate way to infer spatial predications from volume data. However, in order to construct a conic projection we need to determine the geometric position of the vertex. Whereas it is rather intuitive that the vertex should be positioned perpendicularly on the respective anatomical plane (i.e., in X, Y, or Z position) straight above the centroid of the projection of the query structure to this plane, its distance from the query structure might be subject to debate. Spontaneously, one could suggest the body surface. But this is problematic. ▶Figure 5 shows a schematic cross section of the human thorax with the sternum as query structure Q and the two scapulas S1 and S2. In ▶Figure 5a O marks the observation point at the body surface. From this point, even the scapulas would be located posterior to the sternum. This is not plausible. Additionally, the size and the shape of the query volume would greatly vary with the distance between skin and bone, which depends, e.g. on the layer of body fat in between. The query volume that derives from the observation point O′ looks much more plausible. With a distance of about one meter it corresponds to a standard distance between a physician and a patient, or between an anatomist and a cadaver, which represents quite well the context in which terms like "anterior" or "posterior" had been coined and used in a pre-radiology era.

The relations between locative expressions, human cognition, and reality have been subject to a large number of theoretical deliberations and experimental studies. A. Herskovits [35] concluded that simple relation models for locative expressions are mostly inadequate, and she provides numerous examples to demonstrate how the prototypical meaning deviates from the pragmatics of use. Statements like "The North Star is to the left of the mountain peak" suggest the representation of the objects on reality as a geometric scene. This come close to what we have proposed in relation to ▶Figure 4b, where the visual representation of objects is taken into account. Klaus-Peter Gapp [36] conducted experiments on the interdependencies between angle, distance, and shape with respect to the acceptability of projective relations, using the relation "above". His study provided empirical evidence about the importance of the angle between the query and the target object, as postulated in our discussion of ▶Figure 4b. Regier and Carlson [37] developed a more complex computational model in which the size and the shape of the query object are used as additional parameters to predict the plausibility of the locative statements regarding a target object and found a good fitting with experimental results.

According to our analysis, parallel projection volumes for querying volume data using locative expressions are of limited cognitive validity. A new version of Puget et al.'s spatial-symbolic query engine for anatomy should therefore take into account the considerable body of knowledge on spatial perception and related language expressions. A set of experiments using anatomical data and involving medical experts should be conducted in order to optimize the shape of projection volumes, thus increasing the utility of systems for querying and navigating in anatomical volume data.

References

1. Puget A, Mejino LV, Detwiler LT, Franklin JD, Brinkley JF. Spatial-symbolic Query Engine in Anatomy. Methods Inf Med 2012; 51 (6): 463–478.
2. Handels H, Horsch A, Meinzer HP. Advances in Medical Image Computing. Methods Inf Med 2007; 46 (3): 251–253.
3. Handels H, Ehrhardt J. Medical Image Computing for Computer Supported Diagnostics and Therapy – Advances and Perspectives. Methods Inf Med 2009; 48 (1): 11–17.
4. Preim B, Bartz D. Visualization in Medicine. Morgan Kaufmann; 2007.
5. Forkert N, Schmidt-Richberg A, Fiehler J, Illies T, Möller D, Handels H, Säring D. Fuzzy-based Vascular Structure Enhancement in Time-of-Flight MRA Images for Improved Cerebrovascular Segmentation. Methods Inf Med 2011; 50 (1): 74–78.
6. Marschollek M, Teistler M, Bott OJ, Stuermer KM, Pretschner DP, Dresing K. Pre-operative Dynamic Interactive Exploration of Complex Articular Fractures Using a Novel 3D Navigation Tool. Methods Inf Med 2006; 45 (4): 384–388.
7. Handels H, Werner R, Schmidt R, Frenzel T, Lu W, Low DA, Ehrhardt J. 4D Medical Image Computing and Visualization of Lung Tumor Mobility in Spatio-temporal CT Image Data. Int J Med Inform 2007; 76 (3): 433–439.
8. Höhne K, Pflesser B, Pommert A, Riemer M, Schubert R, Schiemann T, Tiede U, Schuhmacher U. A Realistic Model of Human Structure from the Visible Human Data. Methods Inf Med 2001; 40 (2): 83–89.
9. Park JS, Jung Y, Shin DS, Chung MS, Riemer M, Handels H. Generating Useful Images for Medical Applications from the Visible Korean Human. Comp Meth Prog Biomed 2008; 92: 257–266.
10. Ehrhardt J, Handels H, Plötz W, Pöppl SJ. Atlas-based Recognition of Anatomical Structures and Landmarks and Automatic Computation of Orthopaedic Parameters to Support the Virtual Planning of Hip Operations. Methods Inf Med 2004; 43: 391–397.
11. von Berg J, Lorenz C. A Geometric Model of the Beating Heart. Meth Inf Med. 2007; 46 (3): 282–286
12. Heimann T, Wolf I, Meinzer HP. Automatic Generation of 3D Statistical Shape Models with Optimal Landmark Distributions. Methods Inf Med 2007; 46 (3): 282–286
13. Hufnagel H, Ehrhardt J, Pennec X, Ayache N, Handels H, Computation of a Probabilistic Statistical Shape Model in a Maximum-a-posteriori Framework. Methods Inf Med 2009; 48 (4): 314–319.
14. Hacker S, Handels H. A Framework for Representation and Visualization of 3D Shape Variations of Organs in an Interactive Anatomical Atlas. Methods Inf Med 2009; 48 (3): 272–281.
15. Ackerman MJ. The Visible Human Project: A Resource for Education. Acad Med 1999; 74 (6): 667–670.
16. Benson N, Whipple M, Kalet IJ. A Markov Model Approach to Predicting Regional Tumor Spread in the Lymphatic System of the Head and Neck. Proceedings of the American Medical Informatics Association Fall Symposium. In: David W. Bates, John H. Holmes and Gilad Kuperman (eds). American Medical Informatics Association; 2006. pp 31–35.
17. Paluszynski W, Kalet I.. Design Optimization Using Dynamic Evaluation. Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, August 1989, Detroit, Michigan, American Association for Artificial Intelligence; 1989.
18. Kalet IJ, Paluszynski W. Knowledge-based Computer Systems for Radiotherapy Planning. American Journal of Clinical Oncology. 1990; 13: 344–351.
19. Ketting C, Austin-Seymour M, Kalet I, Jacky J, Kromhout-Schiro S, Hummel S, Unger J, Fagan L, Griffin T. Automated Planning Target Volume Generation: An Evaluation Pitting a Computer-based Tool Against Human Experts. International Journal of Radiation Oncology, Biology and Physics. 1997; 37 (3): 697–704.

20. Chang S-K Jungert E. Symbolic Projection for Image Information Retrieval and Spatial Reasoning; Academic Press, San Diego; 1996.
21. Rosse C Mejino JLV. A reference ontology for bioinformatics: the Foundational Model of Anatomy. J Biomed Inform 2003; 36 (6): 478–500.
22. DARPA. The Virtual Soldier Project. http://www.virtualsoldier.us/
23. Harrison SJ Feldman J Perceptual comparison of features within and between objects: a new look. Vision Research 2009; 49: 2790–2799.
24. PICITUP. http://www2.picitup.com/
25. Google Goggles. www.goggles.com
26. Track Id. https://play.google.com/store/apps/details?id=com.sonyericsson.trackid&hl=en
27. W3C – World Wide Web Consortium. http://www.w3.org/TR/
28. Canada Health Infoway. Canada's next generation of health care at a glance https://www.infoway-inforoute.ca/index.php/resources/infoway-corporate/vision/doc_download/81–2015-canada-s-next-generation-of-health-care-summary
29. Australian Health Ministers' Conference. National eHealth Startegy Autralia, 2008. http://www.nehta.gov.au/component/docman/doc_download/626-national-e-health-strategy-summary-dec08
30. Department of Health. The power of information: Putting all of us in control of the health and care information we need, 2012. http://www.dh.gov.uk/prod_consum_dh/groups/dh_digitalassets/@dh/@en/documents/digitalasset/dh_134205.pdf
31. WHO-ITU. National eHealth Strategy Toolkit, 2012. http://imianews.wordpress.com/2012/07/16/national-ehealth-strategy-toolkit-from-who-itu/
32. Piette JD, Lun KC, Moura LA Jr, Fraser HSF, Mechael PN, Powell J, Khojag SR. Impacts of e-health on the outcomes of care in low- and middle-income countries: where do we go from here? Bull World Health Organ 2012; 90: 365–372. doi:10.2471/BLT.11.099069
33. Schulz S, Stenzhorn H, Boeker M, Smith B. Strengths and limitations of formal ontologies in the biomedical domain. RECIIS – Electronic Journal in Communication, Information and Innovation in Health, 2009; 3 (1): 31–45. http://dx.doi.org/10.3395/reciis.v3i1.241en
34. Smith B, Ceusters W, Klagges B, Köhler J, Kumar A, Lomax J, Mungall C, Neuhaus F, Rector AL, Rosse C. Relations in biomedical ontologies. Genome Biol 2005; 6 (5):R46. Epub 2005 Apr 28.
35. Herskovits A. Semantics and Pragmatics of Locative Expressions.Cognitive Science. 1985; 9: 341–378.
36. Gapp K-P. Angle, Distance, Shape, and their Relationship to Projective Relations. In: Moore JD, Lehman JF (eds). Proc of the 17th Annual Conference of the Cognitive Science Society, Mahwah, NJ: Lawrence Erlbaum; 1995. pp 112–117.
37. Regier T, Carlson LA. Grounding spatial language in perception: An empirical and computational investigation. Journal of Experimental Psychology: General 2001; 2 (130): 273–298.

Addresses of the Authors

Wolf-Tilo Balke
Technische Universität Braunschweig
Institute for Information Systems
Mühlenpfordtstraße 23
D-38106 Braunschweig
Germany
E-mail: balke@ifis.cs.tu-bs.de

Heinz Handels
Universität zu Lübeck
Institute of Medical Informatics
Ratzeburger Allee 160
23538 Lübeck
Germany
E-mail: handels@imi.uni-luebeck.de

Ira J. Kalet
University of Washington
Radiation Oncology Department
Box 356043
Seattle, Washington 98195–6043
USA
E-mail: ikalet@uw.edu

Michio Kimura
Hamamatsu University School of Medicine
Medical Informatics Department
1–20–1 Handayama
Hamamatsu
431–3192
Japan
E-Mail: kimura@mi.hama-med.ac.jp

Casimir A. Kulikowski
Rutgers University
Department of Computer Science
Hill Center – Busch Campus
Piscataway, NJ 08854
USA
E-mail: kulikows@cs.rutgers.edu

Lincoln A. Moura
Assis Moura eHealth
Rua Miguel Tostes, 230
90430–060 Porto Alegre, RS
Brazil
E-mail: lincoln.a.moura@gmail.com

Andreas Pommert
Voxel-Man Group
Department of Anatomy and Experimental Morphology
University Medical Center Hamburg-Eppendorf
Martinistraße 52
20251 Hamburg
Germany
E-mail: pommert@uke.de

Stefan Schulz
Medical University of Graz
Institute for Medical Informatics, Statistics and Documentation
Auenbruggerplatz 2
8036 Graz
Austria
E-Mail: stefan.schulz@medunigraz.at