

The Design of an Ontology-Enhanced Anatomy Labeler

Ifeoma Nwogu¹, Jason J. Corso¹ and Thomas Bittner²

¹ Department of Computer Science and Engineering

² Departments of Philosophy and Geography

State University of New York at Buffalo

Buffalo, New York 14260

CSE Technical Report

May 27, 2008

Abstract

In this paper, we present a formal theory for symbolically modeling the spatial relationships that exist between gross-level anatomical structures in the human body. We develop these theories with the goal of computer-based inference. The formal theories are used for building models, which can be applied on graph representations of medical images. We describe an end-to-end design for inferring (labeling) parts of the anatomy from an unlabeled three-dimensional data set. Given a finite set of labels (corresponding to anatomical structures from a taxonomy), we probabilistically assign one label to each node in an anatomical graph. The paradigm we present is generalizable and can help bridge the gap between purely informational ontologies and ontologies for intelligent agents.

1 Motivation and Background

The data from most current medical imaging modalities such as Computerized Tomography (CT), ultrasound, Positive Emission Tomography (PET) etc., are stored in digital format, yet image reading/analysis is still a very manual process with little or no computer intelligence involved. Medical image reading by humans can be a slow and tedious process, prone to human errors while lacking objective means of quantifying and assessing the reader's skills. To overcome this problem in clinical settings, multiple readers are required to review the same data sets, to compensate for any errors that might arise from one reader alone.

We propose an ontology-enhanced probabilistic reasoning system for labeling gross anatomical structures i.e., maximally connected anatomical parts of non-microscopic size (thus anatomical structures like cells and molecules are not gross-level anatomical parts) and narrow our area of focus to CT data in the chest region.

In observing how human readers analyze CT scans to determine if a pathology exists or not, we discovered that they primarily apply (1) their knowledge of the appearances of the gross anatomical structures (such as intensity, shape, size etc.); (2) their knowledge of the locations of structures relative to each other and (3) flag any unusual structures at unexpected locations. Hence, we incorporated these concepts in the design of the anatomy labeler. Different models for representing the appearance of anatomical structures have been investigated extensively in the areas of computer vision and medical image analysis and some popular intensity-based representations include edge/gradient measurements, textures and histograms. Others include spectral features, geometric representations. The study of the location and the arrangement of

anatomical structures within the human body (for interpretation by computers) has been much less investigated and we argue that a good machine-based human anatomy segmentation/labeling system will need to incorporate both concepts.

The representation of spatial relations between body parts is the main component of an anatomical ontology. The quality of such an ontology is determined by how unambiguous the use of its relational terms are, and ambiguity is avoided by a clear understanding/definition of what exactly the relational terms denote.

Foundational Model of Anatomy (FMA) is currently the most comprehensive ontology of human canonical anatomy consisting of approximately 75,000 classes and over 120,000 terms. It has over 2.1 million relationship instances from 168 relationship types and contains organs, their parts, as well as organ systems and body parts (body regions). FMA also includes definitions of macromolecules, cells and their parts, tissues etc. Although much research has gone into the development of the FMA, it is still not strongly suitable for supporting machine-based intelligence in its current state and we discuss this in more detail in the ensuing sections.

The goal of this research work is to provide succinct, unambiguous definitions for the relational terms in an anatomical ontology, integrating the ontology into a probabilistic reasoning framework to better label the anatomical parts of a section of the human body.

2 Symbolic Modeling Principles

The representation of spatial relations between body parts includes mereological relations such as parthood and overlap, topological relations such as connectedness and one-pieceness, as well as location relations such as located in and partially coincides. We use qualitative relations (rather than exact measurements) because of the many variations that often occur in the human anatomy from one individual to another. For example, according to the Atlas of Vascular Anatomy [20] there is a number of variations in the origin of aortic arch branches. There is a usual pattern that occurs in about 65% of the population, a less common pattern in about 27%, and three other varieties, the least popular occurring in only 1.2%. Hence, for an automated reasoning system, it is important to design a universal model where any variations observed will be flagged either as a normal aberration or a pathology.

The qualitative spatial relations will be later exploited for probabilistic reasoning. In this section, we briefly review the qualitative spatial relations as relevant for this paper. For Details see [8].

Parthood, symbolized as P , is the relation that holds between two entities, x and y , whenever x is part of y . For example, the heart is part of the chest (or thorax). Other relations, such as overlap (having a common part) (D_O) are defined in terms of parthood. For example, my hand is part of my body and my hand and my body overlap.

Location relations are defined terms of a region function, r , that maps each spatial entity to the spatial region at which is exactly located at the given moment. The function r is treated as a time-independent primitive hence the notion of movement is absent in our ontology. x is located in y if x 's spatial region is part of y 's spatial region (D_{LocIn}); x and y partially coincide if x 's spatial region and y 's spatial region overlap (D_{PCoin}).

For example, the right lung is located in (but not part of) the thoracic cavity. The esophagus partially coincides with the mediastinal space. The relation $LocIn$ does not hold since the esophagus' region is not part of the region of the mediastinal space, i.e. part of the esophagus lies outside of the mediastinal space. Also, a catheter or stent inside a patients chest is located-in (but not part of) the patients chest. The "located-in" relation does not directly correspond to the English expression "is-in".

Two entities x and y are connected, Cxy , if and only if x and y overlap or x and y are in direct external contact. We define: x and y are externally connected if and only if x and y are connected and x and y do not partially coincide (D_{EC}); x and y are separated if and only if x and y are not connected (D_{SP}).

label	Axiom/Definition
(P1)	Pxx
(P2)	$Pxy \wedge Pyx \rightarrow x = y$
(P3)	$Pxy \wedge Pyz \rightarrow Pxz$
(D _O)	$Oxy \equiv (\exists z)(Pzx \wedge Pzy)$
(L1)	$Pxy \rightarrow Pr(x)r(y)$
(L2)	$r(r(x)) = r(x)$
(D _{LocIn})	$LocIn(x, y) \equiv Pr(x)r(y)$
(D _{PCoin})	$PCoin(x, y) \equiv Or(x)r(y)$
(C1)	Cxx
(C2)	$Cxy \rightarrow Cyx$
(C3)	$LocIn(x, y) \rightarrow (\forall z)(Czx \rightarrow Czy)$
(D _{EC})	$ECxy \equiv Cxy \wedge \neq PCoin(x, y)$
(D _{SP})	$SPxy \equiv \neq Cxy$

Table 1: Axioms and definitions of qualitative spatial relations

In addition to spatial relations between individual entities we define our representations of spatial relations on universal types. An example of an over-arching type would be Liver while its instances or individuals would be Mrs Smith’s heart, Ifeoma’s lungs etc. Between individual entities and their universal types the instantiation relation *Inst* holds. Type-based qualitative relations enable us to establish the connection to canonical anatomy as represented in the FMA [14].

Let R be any binary relation on individuals – for example, the parthood relation (P), the overlap relation (O), the located in relation ($LocIn$), or any of the other relations introduced above. We can use R and the instantiation relation to define the following three relations among types. (See also [7, 4] where these distinctions are made for different versions of type-level parthood relations. [15] uses description logic for distinguishing versions of type-level parthood relations. [18] also define additional type-level relations like adjacent-to, participates-in, etc.)

(Type Relation Definition Schema 1)

$$R_1(A, B) \equiv (\forall x)(Inst(x, A) \rightarrow (\exists y)(Inst(y, B) \wedge Rxy))$$

(Type Relation Definition Schema 2)

$$R_2(A, B) \equiv (\forall y)(Inst(y, B) \rightarrow (\exists x)(Inst(x, A) \wedge Rxy))$$

(Type Relation Definition Schema 1-2)

$$R_{12}(A, B) \equiv R_1(A, B) \wedge R_2(A, B)$$

As [7] puts it, R_1 type-level relations place restrictions on all instances of the first argument. $R_1(A, B)$ tells us that something is true of all A ’s – each A stands in the R relation to some B . R_2 type-level relations place restrictions on all instances of the second argument. $R_2(A, B)$ tells us that something is true of all B ’s – for each B there is some A that stands in the R relation to it. R_{12} type-level relations place restrictions on all instances of both arguments. $R_{12}(A, B)$ tells us that something is true of all A ’s and something else is true of all B ’s– each A stands in the R relation to some B and for each B there is some A that stands in the R relation to it.

A universal type should have at least one instance or entity which, at each moment of its existence, occupies a unique spatial location. The universal related to the entity can be material (lung, muscle) or immaterial (thoracic cavity, mediastinum). Material types have instances with a positive mass while immaterial type instances have no mass.

2.1 Convexity

With only the parthood relations, the connectedness relations, and the region function we cannot introduce the kind of location relation that holds between say, the pleural space and the pleural membrane. The region of my pleural space does not overlap the region of the pleural membrane. The region of the pleural space does not overlap the region of the pleural membrane, rather, the region of the pleural space lies within a region which is somewhat bigger than the region of the pleural membrane, i.e. the convex hull the region of the pleural membrane.

An region can be said to be *convex* if for every pair of points within the region, every point on the straight line segment that joins them is also within the region. The region occupied by a solid cube is convex, but any region that has a dent in it is not convex. Regions occupied by a flower vase or the pleural membrane are not convex. The convex hull of an entity x is the smallest convex region of which x 's region is part. For example, the convex hull of the pleural membrane extends over both the pleural membrane and the space inside the pleural membrane.

We therefore define a new primitive, the convex hull function, ch , which maps every entity to its convex hull. The following axioms capture important properties of the convex hull:

$$(CH1) Pr(x)ch(x) \\ \text{(x's region is part of x's convex hull)}$$

$$(CH2) LocIn(x, y) \rightarrow Pch(x)ch(y) \\ \text{((if x is located in y, then x's convex hull is part of y's convex hull)}$$

$$(CH3) ch(ch(x)) = ch(x) \\ \text{(the convex hull of x's convex hull is x's convex hull)}$$

Using the convex hull function we now define further spatial relations: x is surrounded by y if and only if x 's region is part of y 's convex hull and x 's region does not overlap y 's region (D_{SurrBy}); x is partly surrounded by y if and only if x 's region overlaps y 's convex hull and x 's region does not overlap y 's region ($D_{PSurrBy}$); Symbolically:

$$(D_{SurrBy}) SurrBy(x, y) \equiv Pr(x)ch(y) \wedge \neg Or(x)r(y) \\ \text{(x is surrounded by y)}$$

$$(D_{PSurrBy}) PSurrBy(x, y) \equiv Or(x)ch(y) \wedge \neg Or(x)r(y) \\ \text{(x is partially surrounded by y)}$$

2.2 Summary

The primitive relations and definitions presented to this point are sufficient to characterize the important properties required for performing our probabilistic reasoning. As can be observed, all the relations binary across entities. We are especially interested in the binary relations because they will be used to provide *evidence* either for or against the results of our probabilistic reasoning.

As an example, if the system used probabilistic reasoning to generate a *belief* that a particular section of the anatomy was say, the trachea, then the anatomical ontology can either strengthen or disqualify such a belief by examining the entities with binary relations with the trachea. So if the so-called trachea was surrounded by bone, there would be strong evidence that the result of the probabilistic reasoning was wrong! If on the other hand, the so-called trachea was found to be externally-connected to a pulmonary vessel, the evidence that the entity was truly the trachea would be very strong. This type of evidence-based statistical inference will be discussed in more detail in the ensuing section.

3 Machine Realizability

We define machine realizability as the ability of a computer to extract information by processing input data in a manner consistent with human interpretation of the same data. In observing how human readers analyze CT scans to determine if a pathology exists or not, we discovered that the readers primarily apply (1) their knowledge of the appearances of the gross anatomical structures (such as intensity, shape, size etc.); (2) their knowledge of the locations of structures relative to each other and (3) flag any unusual structures at unexpected locations. For the remaining part of the paper, we describe the machine-realizable design of an Anatomy Labeler implementing concepts (1) and (2) above.

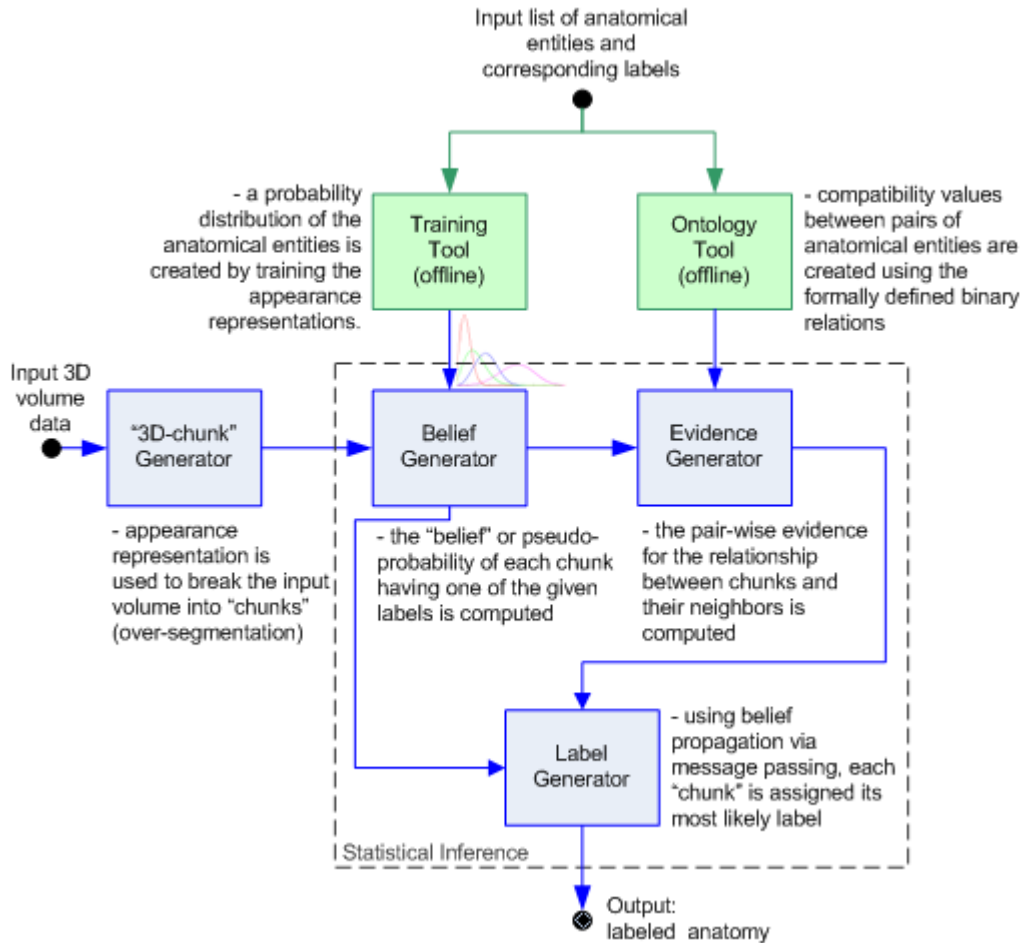


Figure 1: Flowchart of the processes of the ontology-enhanced anatomy labeler

3.1 Overview of the end-to-end process

The first step in the machine realizability process is to identify a suitable form of representing the given image data. The study of the representation of natural images and anatomical structures for statistical inferencing, has been investigated extensively in the areas of computer vision and medical image analysis. Specifically for CT images, most modern equipment has a capacity of 4096 gray tones which represent different density levels in Hounsfield Units (HU). The HU density scale is a linear transformation of the linear attenuation coefficient measurement of substances such that the HU value of water is zero and that of air is

-1000. Therefore, HU in the human anatomy values are typically between -1000 (air) and +1000 (bone-like structures). The HU values alone are used for the appearance representation.

Off-line, a taxonomy (set of controlled vocabulary terms) of the gross anatomical entities-of-interest is created and corresponding labels are assigned (one label for each entity). The taxonomy drives the final output results of the anatomy labeler. The 3D volumes obtained from normal CT scans are manually tagged with the labels corresponding to the different parts of the anatomy. The appearance representations are trained with the labeled data to generate probability distributions of the entities. The anatomical ontology is also created for the items of the taxonomy, paying particular attention to the pairwise relationships. The “learned” distributions and the anatomical ontology provide the back-end statistics to the anatomy labeler.

The input to the labeler is a 3D volume constructed from the CT slices-to-be-investigated. Using the appearance representations, the volume is segmented into smaller “chunks” using the well-tested segmentation-by-weighted-aggregate algorithm (SWA) [5, 16]. The output of the algorithm is a set of local, coherent grouping of voxels which preserves most of the structure necessary for segmentation. Because over-segmentation is an easier problem to solve than true segmentation, we pre-process the data by over-segmenting it.

To infer the most likely label on each chunk, the unary pseudo-probability (called belief since the probabilities do not add up to 1) is computed from the training distributions. The pairwise relationships - parthood, overlaps, located-in, partially-coincides, externally-connected, surrounded-by and partially-surrounded-by are used to compute evidence between any two neighboring chunks. Iteratively updating the beliefs/evidences propagated by its neighbors, every chunk will eventually converge to have maximum evidence of its true label. The process flow is presented pictorially in figure 1.

3.2 Statistical inference from anatomical graphs

Mathematically, the pairwise adjacency relationships between the chunks of the 3D volume can be represented with a graph $G(V, E)$ where V is the set of vertices or ‘nodes’ and E is the set of edges connecting pairs of vertices. A Markov network graphical model is especially useful for explicitly modeling the relationships between the nodes of the graph. We consider the undirected pairwise Markov Random Field (MRF) where each node $i \in V$ has a set of possible states X corresponding to hidden states $x = x_1, \dots, x_n$. If each hidden node is connected to an observed node y_i and $y = \{y_i\}$, then the goal is to infer some information about the states $x_i \in X$. The edges in the graph indicate statistical dependencies. In our anatomical graph labeling problem, the hidden states X are related to the labels identified in the taxonomy. The joint probability function over the entire graph is given by:

$$p(x_1, x_2 \dots x_n | y) = \frac{1}{Z} \prod_i \phi_i(x_i, y_i) \prod_{ij} \psi_{ij}(x_i, x_j) \quad (1)$$

where $\phi_i(x_i, y_i)$ is the local **belief** for node i ; $\psi_{ij}(x_i, x_j)$ is the **compatibility function** between nodes i and j and Z is a normalization constant. We can write $\phi_i(x_i)$ as shorthand for $\phi_i(x_i, y_i)$.

Equation (1) can be solved recursively by making some simplifying assumptions on the graph. Although the recursive solution is obtained for the simple graph (such as a chain), it has been empirically observed to generalize to more complex graphs [21] although convergence is not always guaranteed.

The recursive updating can be mathematically written as:

$$m_{ij}(x_j) \leftarrow k \sum_{x_i} \psi_{ij}(x_i, x_j) \phi_i(x_i) \prod_{k \in N(i) \setminus j} m_{ki}(x_i) \quad (2)$$

$$b_i(x_i) \leftarrow k \phi_i(x_i) \prod_{k \in N(i)} m_{ki}(x_i) \quad (3)$$

where m_{ij} denotes the **message** node i sends to node j , k is a normalization constant, $N(i) \setminus j$ refers to all the neighbors of i , except j and b_i is the approximate marginal posterior probability (or belief) at node i .

Equation(2) describes how messages are propagated across the nodes of G . This message passing algorithm is also known as **belief propagation (BP)**. $\phi_i(x_i)$ is obtained through maximum likelihood estimation using the probability distribution obtained from training. $\psi_{ij}(x_i, x_j)$ is obtained using the formal ontology assigned to entities from the taxonomy.




Figure 2: Illustration of local message passing.

The BP algorithm solves inference on graphical models problems via a series of local message-passing operations. Refer to Figure 2 for an example. To compute the outgoing message (red) in the figure, the central node must combine all incoming messages (blue) with its own inherent belief. In our implementation BP is applied to the graph using discrete variables where the messages are vectors.

3.2.1 Computing initial inherent beliefs via EM-GMM

In order to compute the initial inherent belief ϕ_i , we first manually assign labels to every voxel in the CT data, where each label represents the anatomical structures-of-interest (see Table 2). If there were M labeled tissue classes, by assuming a Gaussian distribution for each tissue label, we learn a mixture of M Gaussian models over the HU values of the labeled data. The resulting observation model is called the Gaussian Mixture Model (GMM).

If $p(\mathbf{x}|\Theta)$ is the density function of the mixture of Gaussian distributions, then the resulting probabilistic model can be written as:

$$p(\mathbf{x}|\Theta) = \sum_{i=1}^M \alpha_i p_i(\mathbf{x}|\theta_i) \quad (4)$$

where we have M component densities mixed together with M mixing coefficients α_i such that $\sum_{i=1}^M \alpha_i = 1$. Also, $\Theta = (\alpha_1, \dots, \alpha_M, \theta_1, \dots, \theta_M)$; each p_i is a Gaussian density function parameterized by $\theta_i = (\mu_i, \sigma_i^2)$.

The inherent belief ϕ_i can then be computed by testing for the maximum likelihood parameters using expectation-maximization (EM) [6].

3.2.2 Computing compatibility functions

The compatibility function $\psi_{i,j}$ relates the chunks or nodes of the graph to each other using the anatomic ontology relationships. We use the simplest interacting Pott's model where the function takes one of 2 values $-1, +1$ and the interactions exists only for amongst neighbors with the following binary relations - parthood(P), overlap(O), located-in($LocIn$), partially-coincides($PCoin$), externally-connected(EC) and

for completeness, $\text{separate}(SP)$, as defined in section 2. Our compatibility function is therefore given as:

$$\psi(x_i, x_j) := \begin{cases} +1 & \text{if } R(i, j) \in \{P, O, LocIn, PCoin, EC\} \\ -1 & \text{otherwise} \end{cases} \quad (5)$$

4 Preliminary Results and Discussion

Table 2 shows some examples of binary relations in our ontology. The level of granularity in the ontology is determined by the tissue structures of the entities. For example, the left lung and the right lung are very different anatomical structures where the left lung typically overlaps the heart and the right one does not, but their tissue structure is exactly the same, hence we define an entity *lung* rather than *right lung*, *left lung*. Similarly, there are many different muscles in the thoracic cavity such as the *deltoid muscle*, *stenothyroid muscle*, *scalenus muscle* etc., but the ontology is defined only with the entity *muscle* at this stage.

The inference algorithm does not consider geometric properties when computing the initial inherent beliefs, it only uses the tissue make-up thus far.

<i>Entity(x)</i>	<i>Relation</i>	<i>Entity(y)</i>
<i>Lung</i>	<i>PSurrBy</i>	<i>Muscle</i>
<i>Lung</i>	<i>EC</i>	<i>Bone</i>
<i>Lung</i>	<i>LocIn</i>	<i>Pleural Cavity</i>
<i>Lung</i>	<i>Overlaps</i>	<i>Heart</i>
<i>Heart</i>	<i>LocIn</i>	<i>Mediastinum</i>
<i>Arteries</i>	<i>LocIn</i>	<i>Mediastinum</i>
<i>Bone</i>	<i>EC</i>	<i>Muscle</i>
<i>Muscle</i>	<i>PSurrBy</i>	<i>Fat</i>
<i>Trachea</i>	<i>SurrBy</i>	<i>Muscle</i>
<i>Esophagus</i>	<i>SurrBy</i>	<i>Muscle</i>
<i>Arteries</i>	<i>PCoin</i>	<i>Heart</i>
<i>Superior mediastinum</i>	<i>LocIn</i>	<i>Thoracic Cavity</i>
<i>Mediastinum</i>	<i>LocIn</i>	<i>Thoracic Cavity</i>
<i>Pleural Cavity</i>	<i>LocIn</i>	<i>Thoracic Cavity</i>

Table 2: Examples of common binary relations of the anatomical ontology of the thoracic region




Figure 3: An example of the “chunking” concept shown on a 2D CT slice. Each chunk becomes a graph node. Adjacent nodes share an edge in the graph.

5 Conclusion

In this paper, we presented the design of an ontologically enhanced method of labeling anatomical structures. To accomplish this, we presented a formal theory for the ontology of modeling anatomical parts, using mereological relations such as parthood and overlap, topological relations such as connectedness and one-pieceness, as well as location relations such as located in and partially coincides. In addition, we presented a probabilistic method of initially labeling the anatomical parts. This involved the generation of a model, via training. Labels were inferred from the model and evidences from the ontology were used to strength the likelihood of the true labels of the structures.

Although this process is currently in the design phase, it is being implemented on 2D slices as a proof-of-concept. Extending to 3D might be challenging because pairwise interactions might not be sufficient for accurate inferences from the graph. But these are issues we intend to handle as the design and implementation progress.

Future work will involve a post-processing step where the level of granularity in the ontology can be raised, to incorporate true anatomical entities such as the *deltoid muscle*, *sternothyroid muscle*, *scalenus muscle* etc., rather than tissue-based entities like *muscle*. Also we intend to investigate incorporating 3D textures for appearance feature representation. Related work in Artificial Intelligence includes [1].

References

- [1] J. Atif, C. Hudelot, G. Fouquier, I. Bloch, and E. Angelini. From generic knowledge to specific reasoning for medical image interpretation using graph-based representations. *Int. Joint Conf. on Artif. Intell. (IJCAI)*, pages 224–229, 2007.
- [2] T. Bittner and L. J. Goldberg. The qualitative and time-dependent character of spatial relations in biomedical ontologies. *Bioinformatics*, doi: 10.1093/bioinformatics/btm155, 2007.
- [3] T. Bittner, L. J. Goldberg, and M. Donnelly. Ontology and qualitative medical images analysis. In Mitsuhiro Okada and Barry Smith, editors, *Proc. InterOntology 08*. Tokyo University Press, 2008.
- [4] Thomas Bittner, Maureen Donnelly, and Barry Smith. Individuals, universals, collections: On the foundational relations of ontology. In *The 3rd Conf. on Formal Ontology in Inf. Sys.*, 2004.
- [5] Jason J. Corso, E. Sharon, and Alan L. Yuille. Multilevel Segmentation and Integrated Bayesian Model Classification with an Application to Brain Tumor Segmentation. In *Proceedings of Medical Image Computing and Computer Aided Intervention (MICCAI)*, volume 2, pages 790–798, 2006.
- [6] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.
- [7] Maureen Donnelly. On parts and holes—the spatial structure of the human body. *Proc. of the 11th World Congress on Medical Informatics (MedInfo-04)*, pages 351–356, 2004.
- [8] Maureen Donnelly and Thomas Bittner. Spatial relations between classes of individuals. In Mitsuhiro Okada and Barry Smith, editors, *Spatial Information Theory*. Springer Berlin / Heidelberg, 2005.
- [9] Maureen A. Donnelly. Formal theory for reasoning about parthood, connection, and location. *Artif. Intell.*, 160:145–172, 2004.
- [10] P. Felzenszwalb and D. Huttenlocher. Efficient belief propagation for early vision. In *Proc. IEEE Conf. Comput. Vision And Pattern Recogn.*, 2004.

- [11] William T. Freeman, Egon C. Pasztor, and Owen T. Carmichael. Learning low-level vision. *International Journal of Computer Vision*, 40(1):25–47, 2000.
- [12] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- [13] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems : Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [14] Cornelius Rosse and Jr. José L. V. Mejino. A reference ontology for biomedical informatics: the foundational model of anatomy. *J. of Biomedical Informatics*, 36(6):478–500, 2003.
- [15] Stefan Schulz, Anand Kumar, and Thomas Bittner. Biomedical ontologies: what part-of is and isn't. volume 39, pages 350–361, 2006.
- [16] E. Sharon, A. Brandt, and R. Basri. Segmentation and boundary detection using multiscale intensity measurements. In *Proc. IEEE Conf. Comput. Vision And Pattern Recogn.*, volume 1, pages 469–476, 2001.
- [17] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [18] Barry Smith, Werner Ceusters, Bert Klagges, Jacob Kohler, Anand Kumar, Jane Lomax, Chris Mungall, Fabian Neuhaus, Alan Rector, and Cornelius Rosse. Relations in biomedical ontologies. *Genome Biology*, 6(5):R46, 2005.
- [19] Richard Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall F. Tappen, and Carsten Rother. A comparative study of energy minimization methods for markov random fields. In *ECCV (2)*, volume 3952 of *Lecture Notes in Computer Science*, pages 16–29, 2006.
- [20] Renan Uflacker. In *Atlas of Vascular Anatomy: An Angiographic Approach*. Lippincott Williams and Wilkins, 2006.
- [21] J. Yedidia, W. Freeman, and Y. Weiss. Bethe free energy, kikuchi approximations, and belief propagation algorithms. MERL Tech. Report., TR2001-16., 2001. <http://www.merl.com/reports/docs/TR2001-16.pdf>.