From Generic Knowledge to Specific Reasoning for Medical Image Interpretation using Graph based Representations

Jamal Atif, Céline Hudelot, Geoffroy Fouquier, Isabelle Bloch, Elsa Angelini

GET - Ecole Nationale Supérieure des Télécommunications - Dept TSI - CNRS UMR 5141 LTCI

46 rue Barrault, 75013 Paris, France – isabelle.bloch@enst.fr

Abstract

In several domains of spatial reasoning, such as medical image interpretation, spatial relations between structures play a crucial role since they are less prone to variability than intrinsic properties of structures. Moreover, they constitute an important part of available knowledge. We show in this paper how this knowledge can be appropriately represented by graphs and fuzzy models of spatial relations, which are integrated in a reasoning process to guide the recognition of individual structures in images. However pathological cases may deviate substantially from generic knowledge. We propose a method to adapt the knowledge representation to take into account the influence of the pathologies on the spatial organization of a set of structures, based on learning procedures. We also propose to adapt the reasoning process, using graph based propagation and updating.

1 Introduction

Several domains, such as anatomy, can benefit from a strongly structured knowledge, that can be expressed for instance by an ontology. This knowledge is of great help for interpreting particular cases or instantiations (e.g. individual anatomy). As a typical application, we consider in this paper the problem of knowledge-based recognition of anatomical structures in 3D brain images. The available knowledge includes generic descriptions of structures, and their spatial organization, in particular their spatial relations, usually formulated in natural language or formalized in medical ontologies. This knowledge involves complex textual descriptions, which are difficult to translate into an operational model assisting image segmentation and recognition algorithms. Towards this aim, we propose to model this type of anatomical knowledge using graphs, as a natural representation of both objects and relations and a powerful reasoning tool (Section 2). Spatial reasoning in medical image interpretation strongly relies on spatial relations, that are less subject to inter-individual variability than properties of anatomical structures such as size or shape, as described in Section 3. A generic graph model is used in a reasoning process guiding the recognition of anatomical structures in brain magnetic resonance images

and is then instantiated in order to account for specificities of the patient. However, specific cases can exhibit significant deviations from the generic model. This is typically the case in medical applications when pathologies (such as brain tumors) may appear as additional objects in the images, which are not represented in the generic model. Another contribution of this paper, detailed in Sections 4 and 5, is the adaptation of the reasoning process to deal with specific cases, according to both generic knowledge about pathologies and their potential impact on normal structures, and specific patient's information derived from the images. We propose to learn the variability of the spatial relations in the presence of a tumor by using a database of pathological cases (Section 4). The adaptation of the reasoning procedure relies on knowledge about the pathologies, through the exploitation of a brain tumor ontology, the learning results and a graph based propagation process which consists in updating the graph structure and attributes based on image segmentation results (Section 5).

The proposed approach, illustrated in Figure 1, contributes to fill the gap between two widely addressed problems: segmentation and recognition of objects in images on the one hand, and knowledge representation on the other hand. Recently, few approaches for AI based image interpretation have been developed exploiting knowledge and generic information [Crevier and Lepage, 1997], in particular structural knowledge represented with graphs, to drive specific recognition procedures (see e.g. [Mangin *et al.*, 1996; Deruyver *et al.*, 2005; Colliot *et al.*, 2006] among others). However, these methods mainly deal with normal cases, and cannot be applied directly to pathological cases. Noticeably, very little work uses ontology-based approaches for automatic image interpretation, while they are more developed for image retrieval [Smeulders *et al.*, 2000]. Unfortunately, no modeling approach of spatial relations in the image domain has been proposed, and segmentation tasks are not addressed. This paper proposes an original method towards these aims.

2 Generic Knowledge Representation

In several domains, scene interpretation benefits from knowledge about its structural organization and its context. This is typically the case in medical image interpretation. As an example, we consider the case of brain imaging. The brain is usually described as a hierarchical organization. Each level
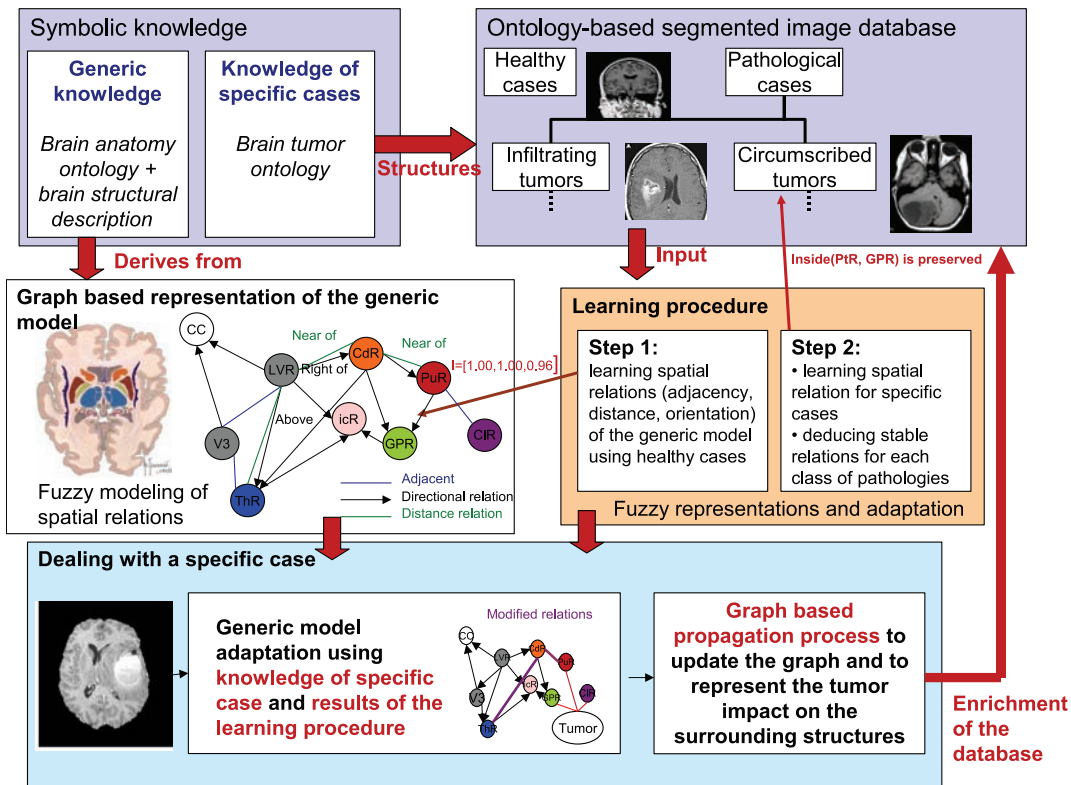
Figure 1: Overview of our framework, illustrating the ontological, graph-based representations, learning and updating procedures. The schematic representation of the generic anatomy is from [Hasboun, 2005].

of the hierarchy is composed of a set of objects, at a given level of granularity. These objects are organized in space in a roughly persistent way. This hierarchical and spatial organization is an important component of linguistic descriptions of anatomical knowledge [Bowden and Martin, 1995; Hasboun, 2005]. Recent developments in the ontology community have shown that ontologies can efficiently encode generic and shared knowledge of a domain. For instance, the Foundational Model of Anatomy (FMA) [Rosse and Mejino, 2003] provides an ontology of the canonical anatomy of the human body.

Based on both linguistic and ontological descriptions, we propose to model the spatial organization of the brain as an hypergraph. Each vertex represents an anatomical structure, while edges or hyperedges carry information about the spatial relations between the vertices they link. The choice of hypergraphs is motivated by the importance of complex relations of cardinality higher than two, such as "between". Moreover this type of structural representation is appropriate for model-based recognition.

Model-based recognition requires a second level of knowledge representation, related to the semantics of the spatial relations in images. Fuzzy representations are appropriate in order to model the intrinsic imprecision of several relations (such as "close to", "behind", etc.), the potential variability (even if it is reduced in normal cases) and the necessary flexi-

bility for spatial reasoning [Bloch, 2005]. Two types of questions are raised when dealing with spatial relations: (i) given two objects (possibly fuzzy), assess the degree to which a relation is satisfied; (ii) given one reference object, define the area of space in which a relation to this reference is satisfied (to some degree). In this paper, we deal mainly with the second question (see Section 3). Therefore we rely on spatial representations of the spatial relations: a fuzzy set in the spatial domain defines a region in which a relation to a given object is satisfied. The membership degree of each point to this fuzzy set corresponds to the satisfaction degree of the relation at this point [Bloch, 2005]. An example is illustrated in Figure 2.


Figure 2: Left: a fuzzy reference object. Right: fuzzy region representing the relation "to the left of" the reference object. Membership degrees vary from 0 (black) to 1 (white).

We now describe the modeling of the main relations that we use: adjacency, distances and directional relative positions.

A **distance** relation can be defined as a fuzzy interval f of trapezoidal shape on $\mathbb{R}^+$, as illustrated in Figure 3. A fuzzy subset μ_d of the image space $\mathcal{S}$ can then be derived by combining f with a distance map d_A to the reference object A:

$$\forall x \in \mathcal{S}, \ \mu_d(x) = f(d_A(x)), \tag{1}$$

where $d_A(x) = \inf_{y \in A} d(x, y)$.

Figure 3: Fuzzy interval representing a distance relation. For instance the relation "close to" can be modeled by choosing $n_1 = n_2 = 0$.

Directional relations are represented using the "fuzzy landscape approach" [Bloch, 1999]. A morphological dilation δ_{ν_α} by a fuzzy structuring element ν_α representing the semantics of the relation "in direction α" is applied to the reference object A: $\mu_\alpha = \delta_{\nu_\alpha}(A)$, where ν_α is defined, for x in $\mathcal{S}$ given in polar coordinates (ρ, θ), as:

$$\nu_\alpha(x) = g(|\theta - \alpha|), \tag{2}$$

where g is a decreasing function from $[0, \pi]$ to $[0, 1]$, and $|\theta - \alpha|$ is defined modulo π. This definition extends to 3D by using two angles to define a direction. The example in Figure 2 has been obtained using this definition.

Adjacency is a relation that is highly sensitive to the segmentation of the objects and whether it is satisfied or not may depend on one point only. Therefore we choose a more flexible definition of adjacency, interpreted as "very close to". It can then be defined as a function of the distance between two sets, leading to a degree of adjacency instead of a Boolean value:

$$\mu_{adj}(A, B) = h(d(A, B)) \tag{3}$$

where $d(A, B)$ denotes the minimal distance between points of A and B: $d(A, B) = \inf_{x \in A, y \in B} d(x, y)$, and h is a decreasing function of d, from $\mathbb{R}^+$ into $[0, 1]$. We assume that $A \cap B = \emptyset$.

In all these definitions, the satisfaction degree of a relation depends on a function (f, g or h) which is chosen as a trapezoidal shape function for the sake of simplicity. A learning step, presented in Section 4, defines the parameters of these functions based on a set of segmented images.

3 Spatial Reasoning for Model-Based Recognition

For the sake of completeness, we summarize in this section previous work on structure recognition using spatial relations and graph based knowledge representations. Details can be found in [Colliot *et al.*, 2006]. The approach is progressive, in the sense that objects are recognized sequentially and

their recognition makes use of knowledge about their relations with respect to other objects. This knowledge is read in the graph representing generic knowledge. The graph also drives the order in which objects are searched. Relations with respect to previously obtained objects have generally different natures and have to be combined in a fusion procedure, at two different levels. First, fusion of spatial relations occurs in the spatial domain, using spatial representations of relations. The result of this fusion allows to build a fuzzy region of interest in which the search of a new object will take place, in a process similar to focalization of attention. In a sequential procedure, the amount of available spatial relations increases with the number of processed objects. Therefore, the recognition of the most difficult structures, usually treated in the last steps, will be focused in a more restricted area. Another fusion level occurs during the final decision step, i.e. segmentation and recognition of a structure. For this purpose, it was suggested in [Colliot *et al.*, 2006] to introduce relations in the evolution scheme of a deformable model, in which they are combined with other types of numerical information, usually edge and regularity constraints. This approach leads to very good results in normal cases.

4 How to Deal with Specific Cases

The method presented so far is not well adapted to cases that greatly differ from the generic model. Particularly, in medical images, the presence of a tumor may induce not only an important alteration of the iconic and morphometric characteristics of its surrounding structures but also a modification of the structural information. We propose a pathology-dependent paradigm based on the segmentation of the pathology and on the use of the extracted information to adapt both the generic graph representation and the reasoning process to specific cases. In this paradigm, ontologies and fuzzy models convey important information to deal with complex situations. In this section, we address the following question: given a pathology, which spatial relations do remain stable, and to which extent?

4.1 A Brain Tumor Ontology

In addition to canonical anatomy, image interpretation in diseased cases can benefit from knowledge on pathologies. For example, brain tumor classification systems are highly used in clinical neurology to drive the selection of a therapeutical treatment. Brain tumors are classified according to their location, the type of the tissue involved, their degree of malignancy and other factors. The main brain tumor classification system is the WHO grading system [Smirniotopoulos, 1999] which classifies brain tumors according to histological features and radiologic-pathologic considerations. Differential diagnosis of brain tumor based on the location of the tumor was also proposed[1]. In this paper, we propose a brain tumor ontology which encodes these different kinds of knowledge. Named tumors (e.g. *Gliomas*, *Astrocytoma*) are hierarchically organized according to the type of the tissue involved in the tumor. Then for each type of tumors, the

---

[1] http://rad.usuhs.mil/rad/location/location_frame.html

ontology describes their possible locations, their spatial behavior (i.e. infiltrating vs. circumscribed), their composition (i.e. solid, cystic, necrotic, with surrounding edema), their modality-based visual appearance and their grade in the WHO grading system, as shown in Figure 4.
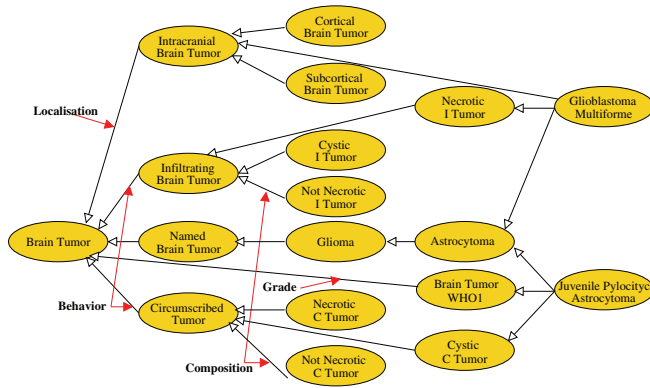

Figure 4: Overview of a subpart of the brain tumor ontology.

This ontology was developed in the framework of Protégé[2] and can be obtained on demand. We show in the next sections how to use this ontology.

4.2 Learning Spatial Relation Stability in Presence of Pathologies

The presence of a pathology can affect the generic structural information in several ways: (1) a spatial relation between two anatomical structures can be invalidated; (2) a spatial relation between two anatomical structures can be validated but with more variability; (3) new relations between anatomical structures and pathological structures can be added; and (4) an anatomical structure can be destroyed by the pathology.

These modifications depend on the type of the relations, on their stability, on the precision or vagueness of the available knowledge. Intuitively, topological relations imply less unstability than metric ones. The nature of the deformation itself (i.e. the nature of the tumor in our case: infiltrating, destroying...) has also an impact on the stability of the spatial relations.

However, these considerations remain intuitive and do not lead to definite conclusions on the nature of the impact. Therefore, we propose to learn the stability of the spatial relations in the presence of tumoral pathologies on a set of examples.

Learning Database

The database is constituted of 18 healthy MRI and 16 pathological MRI where the main anatomical structures were manually segmented. The healthy cases are the images of the widely used IBSR database[3] (real clinical data). The pathological cases include intracranial brain tumors belonging to n different classes. We focus our study on the stability of spatial relations between internal brain structures, namely: ven-

tricles, caudate nuclei, thalami and putamen, which is a clinically relevant choice, according to medical experts' opinion.

The first step concerns the structuration and clustering of the database according to the brain tumor ontology. A key point is the representativity of the database according to predominant spatial behaviors of brain tumors, i.e. their tendency to spread, to destroy (necrotic or not), to stem (cystic tumor, edema presence) and their location.

Learning procedure

From this database, the parameters involved in the construction of the fuzzy representations of spatial relations (for functions f, g, h in particular) are learned. Let us first introduce some notations and definitions: $K = (K^N, K^{P_1}, ... K^{P_n})$ is the learning database with K^N the set of healthy instances and $K^{P_i}, i \in 1...n$, the set of pathological instances of class i. Let c be an instance of K (an image), O_c the set of segmented objects in c and R a spatial relation. We denote by μ_R^N the fuzzy subset in the image space corresponding to the relation R for a healthy case, and by $\mu_R^{P_i}$ the fuzzy subset in the image space corresponding to the relation R for a pathological case of class i.

A leave-one-out procedure is used to learn, for a given spatial relation R, the parameters of its fuzzy formulation μ_R. Since μ_R is defined in the spatial domain, we can directly compare μ_R and the target objects. The parameters are optimized so as to maximize the inclusion of the target object in μ_R (i.e. the object has to fulfill the relation with the highest possible degree). For all $c \in K^k$, $k \in \{N, P_1, ...P_n\}$, $(A_c, B_c) \in O_c$, we compute the fuzzy set μ_R^k with respect to A_c, and for a given inclusion measure $\mathcal{I}$, we compute $\mathcal{I}(B_c, \mu_R^k)$. The fuzzy set optimizing this criterion is denoted by μ_R^c.

In order to learn functions f, g and h (Equations 1, 2 and 3), the minimum or maximum of the values (distances, angles...) are computed for all instances, from which the function parameters are then determined. Let us detail the example of the relation "close to". The training consists in the computation of the maximum distance from a point x of the target object B_c to the reference object A_c:

$$d_{max}^c = \max_{x \in B_c}(d_{A_c}(x)). \tag{4}$$

Then the mean m^k and standard deviation σ^k of the values $\{d_{max}^c\}_{c \in K^k}$ are computed. The fuzzy interval f is then defined as a fuzzy subset of $\mathbb{R}^+$, with kernel $[0, m^k]$ and support $[0, m^k + 2\sigma^k]$. This allows taking into account the variability of the parameters in the training set. An example is illustrated in Figure 5.

A similar approach is applied for adjacency and directional relations.

Stability assessment

The stability of the spatial relations can now be assessed by comparing the learned parameters for specific cases and for healthy ones. A suitable choice for such a comparison is a M-measure of resemblance, according to the classification proposed in [Bouchon-Meunier *et al.*, 1996].

We use a set-theoretic derived M-measure of resemblance defined as the cardinality of the intersection of two fuzzy sets


Figure 5: Learning the relation "close to" between putamen and caudate nucleus μ_d on normal cases and on pathological cases for a class P_i (here high grade gliomas that shift the putamen away from the caudate nucleus).

μ and μ', normalized by the cardinality of their union:

$$\mathcal{R}(\mu, \mu') = \frac{\sum_{d \in \mathcal{D}} \min(\mu(d), \mu'(d))}{\sum_{d \in \mathcal{D}} \max(\mu(d), \mu'(d))}$$

where $\mathcal{D}$ denotes the definition domain of the fuzzy sets. This choice is motivated by the properties of this measure (reflexive, symmetrical, increases with the overlapping between the two fuzzy sets, decreases with their difference).

This resemblance measure is applied on the fuzzy sets learned for each type of spatial relations, as the ones illustrated in Figure 5 (in this case $\mathcal{D}$ is the distance space, i.e. $\mathbb{R}^+$).

Results

In this section, we show some results for an instance of a pathological class (here, high grade glioma, which are spread and destroying). For each spatial relation, we compute the fuzzy resemblance between the learned fuzzy sets in the pathological class and in the healthy one as explained before.

	caudate	ventricle	thalamus	putamen
caudate		1.000	0.713	0.594
ventricle	1.000		0.834	0.416
thalamus	0.573	0.689		0.597
putamen	0.537	0.426	0.653	

Table 1: Degree of resemblance between the fuzzy representations of the the adjacency relation for the healthy class and for a pathological class P_i (high grade gliomas).

	caudate	ventricle	thalamus	putamen
caudate		0.454	0.211	0.194
ventricle	0.639		0.490	0.598
thalamus	0.559	0.537		0.379
putamen	0.530	0.349	0.219	

Table 2: Degree of resemblance for the relation "close to" .

Table 1 shows that the adjacency relation exhibits high resemblance values for structures that have a high degree of adjacency, such as caudate nuclei and ventricles (i.e. the adjacency relations between two structures in a healthy case are similar as those in a pathological case), which confirm our assumption that the adjacency remains stable even in pathological configurations. The spatial relation "close to" is more prone to unstability, in particular for structures that are more affected by the tumor. In the considered pathological class,

the tumor has a strong impact on the putamen for instance. Table 2 shows that the distance relations between the putamen and other structures have a high variability between healthy cases and pathological ones (i.e. low resemblance values). These results are in agreement which visual observations on the images, as illustrated in Figure 6. Note that the learning process in not symmetrical in A_c and B_c and depends on the reference object (Equation 4), which explains that the tables are not symmetrical.


Figure 6: A normal case and a pathological one. The caudate nuclei and the ventricles are adjacent in both cases (hence a high resemblance value for this relation). The putamen is deformed in the pathological case, thus modifying its distance to the caudate nuclei, which explains the low value of the resemblance in Table 2.

5 Knowledge and Reasoning Adaptation for Specific Cases

In this section, we propose an original approach to reason on specific cases based on three steps: (1) detection and segmentation of the tumor from the specific image as described in [Khotanlou *et al.*, 2005]; (2) adaptation of the generic model to the specific case. The idea is to exploit the ontology, enhanced by the learning results, and then add tumor information extracted from images in the graph. Due to the variability of individual cases, this knowledge is only a guide for the recognition of image structures. The real specificity of the processed case is computed by a graph based propagation segmentation process which evaluates the tumor impact on surrounding structures; (3) updating of the database using this new processed case.

5.1 Generic Knowledge Adaptation

The knowledge adaptation process depends on the available knowledge. If we have expert knowledge about the processed case, such as the diagnosis associated with the pathological image, then it can be used to extract its characteristics from the brain tumor ontology and then use the learned relations corresponding to its class in the segmentation process. If we do not have any information about the current case, we can use image matching procedures through a similarity measure to find the best matching class in the database.

The knowledge adaptation process is based on the modification of the generic graph: we first segment the brain tumor, add a tumor node in the graph and localize it in the brain, based on segmentation results; we then modify edge attributes with the learned spatial relation parameters for the class of this particular tumor.

5.2 A Graph Based Propagation Process

The aim of the graph based process is to quantify the real impact of the pathology on surrounding structures by using a progressive model based recognition approach as described in Section 3. The model used is the one which results from the knowledge adaptation step described in the previous section. Starting from a reference structure (usually the lateral ventricle), we propagate the tumor impact in the graph until a behavior which is similar to the healthy case is reached, i.e. until structures which are not affected by the tumor.

The knowledge is read in the graph representing the specific learned knowledge for this class. The graph drives the order in which objects are searched but contrary to the generic case, the choice of the structure to segment is also driven by the stability of the learned relations. More precisely, starting from a previously recognized object A_c, all objects B_c^j linked with an edge to A_c are potential candidates for the next object to be found. The choice of one of these is based on the maximum of resemblance. Usually several relations are shared by A_c and B_c^j. The maximum is then defined based on a lexicographic order on the relations: 1. adjacency, 2. direction, 3. distance. Then we segment the corresponding structure in a similar way as in Section 3. The only difference is that each relation is taken into account with a weight given by the resemblance degree. At last we update segmented structure properties and relations with known structures in the specific graph by edge and node attribute modifications. The process stops when the computed attributes do not widely deviate from the generic model (hence no more updating is required).

An illustrative example of segmentation results obtained using the proposed approach on a pathological case is shown in Figure 7. On this example, once the tumor is segmented, the procedure leads to sucessful segmentation of ventricles, caudate nuclei, thalamus and putamen (in this order). Note that these results would have been very difficult to obtain based on the generic knowledge only.
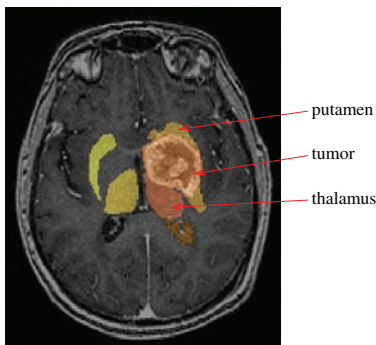

Figure 7: An axial slice of a 3D MRI, with segmented tumor and some anatomical structures.

6 Conclusion

This paper is a contribution to AI based methods for image interpretation. An important feature of the proposed approach is the adaptation of generic knowledge to specific cases taking into account information extracted from individual images, as illustrated in the domain of pathological brain imaging. The method combines in an original way graph representations, fuzzy models of spatial relations, learning procedures and graph propagation. Further work aims at developing the last step, which consists in updating the database with the new processed cases. This includes adapting the shape of the membership functions as in the learning step. Evaluation of the proposed models and the whole approach by medical experts will also be performed.

References

[Bloch, 1999] I. Bloch. Fuzzy Relative Position between Objects in Image Processing: a Morphological Approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(7):657–664, 1999.

[Bloch, 2005] I. Bloch. Fuzzy Spatial Relationships for Image Processing and Interpretation: A Review. *Image and Vision Computing*, 23(2):89–110, 2005.

[Bouchon-Meunier *et al.*, 1996] B. Bouchon-Meunier, M. Rifqi, and S. Bothorel. Towards general measures of comparison of objects. *Fuzzy Sets and Systems*, 84(2):143–153, 1996.

[Bowden and Martin, 1995] D. M. Bowden and R. F. Martin. Neuronames brain hierarchy. *NeuroImage*, 2:63–83, 1995.

[Colliot *et al.*, 2006] O. Colliot, O. Camara, and I. Bloch. Integration of Fuzzy Spatial Relations in Deformable Models - Application to Brain MRI Segmentation. *Pattern Recognition*, 39:1401–1414, 2006.

[Crevier and Lepage, 1997] D. Crevier and R. Lepage. Knowledge-based image understanding systems: a survey. *Computer Vision and Image Understanding*, 67(2):160–185, 1997.

[Deruyver *et al.*, 2005] A. Deruyver, Y. Hodé, E. Leammer, and J.-M. Jolion. Adaptive pyramid and semantic graph: Knowledge driven segmentation. In *Graph-Based Representations in Pattern Recognition: 5th IAPR International Workshop*, volume 3434 / 2005, pages 213–223, Poitiers, France, 2005.

[Hasboun, 2005] D. Hasboun. *Neuranat*. http://www.chups.jussieu.fr/ext/neuranat/index.html, 2005.

[Khotanlou *et al.*, 2005] H. Khotanlou, J. Atif, O. Colliot, and I. Bloch. 3D Brain Tumor Segmentation Using Fuzzy Classification and Deformable Models. In *WILF*, volume LNCS 3849, pages 312–318, Crema, Italy, sep 2005.

[Mangin *et al.*, 1996] J.-F. Mangin, V. Frouin, J. Régis, I. Bloch, P. Belin, and Y. Samson. Towards Better Management of Cortical Anatomy in Multi-Modal Multi-Individual Brain Studies. *Physica Medica*, XII:103–107, 1996.

[Rosse and Mejino, 2003] C. Rosse and J. L. V. Mejino. A Reference Ontology for Bioinformatics: The Foundational Model of Anatomy. *Journal of Biomedical Informatics*, 36:478–500, 2003.

[Smeulders *et al.*, 2000] A.W.M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12):1349–1380, 2000.

[Smirniotopoulos, 1999] J.G. Smirniotopoulos. The new WHO classification of brain tumors. *Neuroimaging Clin N Am*, 9(4):595–613, 1999.