Knowledge Acquisition for Image Understanding

Tatjana Zrimec†* and Claude Sammut*

†Faculty of Computer and Information Science
University of Ljubljana
Trzaska 25, 1001 Ljubljana, Slovenia

tatjana@fri.uni-lj.si
Tel: +386-61-1768-390
Fax: +386-61-1264-647

*School of Computer Science and Engineering
University of New South Wales, Sydney, Australia 2052

claude@cse.unsw.edu.au
Tel: +61-2-9385-3933
Fax: +61-2-9385-1814

Abstract

Image understanding often requires extensive background knowledge. For example, to the lay person, an x-ray image of the human cerebral vasculature can be almost meaningless. However, a trained radiologist, with a knowledge of anatomy and years of experience, can usually quickly detect anomalies in the blood vessels in the x-ray image. The problem addressed in this paper is how the background knowledge for a complex image processing task can be acquired. We discuss how relational machine learning methods can be used to automatically build rules for classifying abnormalities and we speculate on how interactive knowledge acquisition tools, such as ripple-down rules, may be used to refine to the knowledge.

1 Introduction

Image understanding often requires extensive background knowledge. For example, to the lay person, an x-ray image of the human cerebral vasculature can be almost meaningless. However, a trained radiologist, with a knowledge of anatomy and years of experience, can usually quickly detect anomalies in the blood vessels in the x-ray image. The problem addressed in this paper is how the background knowledge for a complex image processing task can be acquired.

Everyone is different and there is considerable variation in the anatomy of different patients. However, the main structures of, say, the blood vessels in the brain must be present. Therefore, general text-book knowledge of anatomy and atlases of the brain can serve as a starting point to construct a framework for the knowledge needed to interpret x-ray angiograms[1] of the brain. But how can variations in individuals be accounted for?

There are text books on variations of human anatomy. These are large volumes that contain descriptions of cases encountered by physicians over many years. In addition, individual radiologists will accumulate their own "case book" over many years of examining patients. Thus, the ability to interpret x-ray images requires a substantial amount of experience beyond simply learning "normal" anatomy. Furthermore, a radiologist's expertise usually includes "procedural knowledge" which guides the interpretation process.

In many meetings with radiologists, they were asked to perform an image interpretation while providing a verbal commentary. We found that the radiologists use different kinds of knowledge which gives them various clues during an investigation and interpretation of images. Expert knowledge includes the following:

- symmetry (displacement of normal structures from the midline),
- deviation from normal patterns (anatomically correct but with some topological variations),
- knowledge of the typical patterns associated with a particular disease,
- abnormal patterns which may not belong to a recognised category.

Radiologists can easily detect patterns in angiograms that show a particular disease. For example, a narrowing of a blood vessel indicates a stenosis. A stroke is manifested by blocked vessels. In an Arteriovenous Malformation (AVM), enlarged vessels and early filling of some veins can be detected. In the case of a tumour a few patterns can be observed: vessel displacement, vessels pushed away and the Sylvian point is depressed. In the case of an AVM, radiologists also use signs for recognising a feeding vessel such as: subtle enlargements, increased flow and changing contrast at the entrance or vessel disappearing.

In all these cases, it was easier to ask a radiologist to demonstrate the problem in an image and what to expect than to try to formalise their descriptions. Further, using examples of cases, they communicated more knowledge than if they were asked to

---

[1] Angiograms are obtained by injecting a dye into a blood vessel and tracing the path of the dye through successive x-ray images.

provide general information. Thus, it is only possible to acquire such domain specific knowledge by learning from images which have been "processed" by experts.

In the following sections of this paper, we will describe a framework for representing anatomical and procedural knowledge for interpreting x-ray angiograms of the human cerebral vasculature. We will also discuss how relational machine learning methods can be used to automatically build rules for classifying types of blood vessels and finally we will speculate on how interactive knowledge acquisition tools, such as ripple-down rules, may be used to refine the knowledge.

2 Modelling the Human Cerebral Vasculature

Apart from the radiologist's expertise, we have four main sources of information:

- a symbolic description of the human anatomy from textbooks (Nolte 1993);
- structural information of a human brain, depicted in the form of an atlas (Talairach 1988);
- a 3D volumetric representation of the blood vessels of the patient's brain, from a set of slices of Magnetic Resonance Angiograms (MRA);
- several x-ray angiograms of the patient's brain in a form of 2D projections.

Models to assist image processing must include not only knowledge of anatomy but also how that anatomy should appear in x-ray images taken from different angles. Thus, a representation of the information from the above sources will contain a symbolic description of the anatomy and the structure and geometry of the vasculature of the brain. It is also necessary to store information about spatial relationships among the vessels and other brain features and how they appear from different views.

We have chosen a frame system to implement the model (Minsky 1975, Goldstein 1977, Horn 1992). The central idea of frames is that knowledge can be stored in a library of frames, which are packets of knowledge that provide descriptions of typical objects and events. Objects are represented by a set of attributes (also called slots) and their associated values. *Generic* frames describe the properties of an entire class of objects, while *instance* frames describe individual objects. An *isa* slot indicates that an instance is a member of a particular class. An *ako* (a kind of) slot indicates that a generic frame is a subclass of another class. Thus, frame systems provide inheritance mechanisms common to most object-oriented systems. *Demons*, are procedures attached to a slot. Specialised demons are triggered automatically, depending on how the slot is accessed. For example, when a new value is added to a slot, the "if added" demon is automatically invoked by the frame system.

An example of the frame representation is shown in Figure 1. This contains a description of the Common Carotid Artery. English expressions are used for purposes of illustration but, in practice, we use the frame system that is incorporated into *iProlog* (Sammut, 1997). *iProlog* is a version of ISO Prolog that has extensions which provide a surface syntax for frame representations. The frames themselves are stored in the usual Prolog manner as relations in the database. This system combines the convenience of a frame language with the pattern matching and deductive capabilities of Prolog. The frame

name:	CCA (Comon Carotid Artery)
ako:	artery
size:	large
width:	5 mm
anterolaterally:	crossed by
	superior thyroid vein
	middle thyroid vein
	jugular vein
laterally:	internal jugular vein
posterolaterally:	vagus nerve
medially:	larynx
	pharynx
	trachea
	esophagus
important features:	embedded in the carotid sheath
begining:	from brachiocephalic location, at the right sternoclavicular joint
termination:	at the upper border of the thyroid cartilage below the angle of the mandible
branches:	ICC (internal common carotid)
	ECC(external common carotid)

name:	LCCA (Left Comon Carotid Artery)
ako:	CCA
begining:	from arch of the aorta in the superior mediastinum

name:	RCCA (Right Comon Carotid Artery)
ako:	CCA
begining:	from brachiocephalic artery behind the right sternoclavicular joint

Figure 1. The description above represents an idealised Common Carotid Artery which also inherits all the properties of a generic artery. There are two Common Carotid Arteries in the cerebral vasculature, these are represented by their own generic descriptions. They will inherit all the properties of the Common Carotid Artery because their ako slots states that each is a kind of Common Carotid Artery.

representation was chosen in favour of Prolog's native relational form because the cerebral vasculature is naturally hierarchical structure and therefore maps easily onto frames. We will see later that demons also allow us to conveniently encode the expert's procedural knowledge.

Three frame hierarchies are used to represent:

- brain structure (the lobes of the brain, etc, excluding the vasculature);
- the vasculature;
- the different views provided by the imaging systems.

Illustrated in Figure 2, each hierarchy contains links to the others, showing the relationships between the different structures. For example, a frame representing a brain structure will contain slots indicating which vessels supply and drain that structure. The frames for those vessels have corresponding slots to indicate which structure they supply or drain.

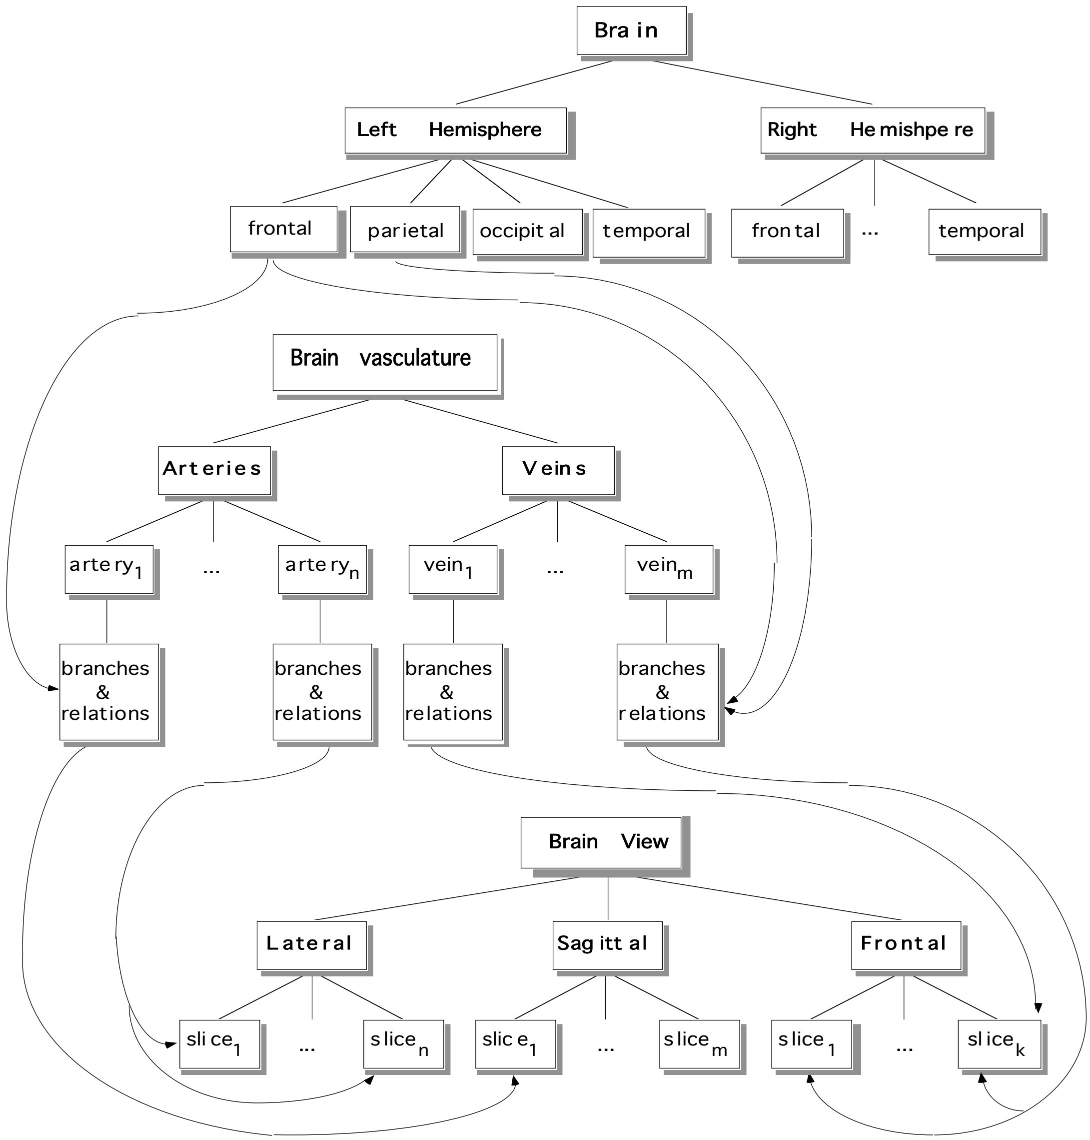

Figure 2. Frame Hierarchy

We will now describe these three frame types in further detail.

2.1 Vasculature Frames

The following slots are used to describe a blood vessel:

- the frame's identifier and the long name of the blood vessel;
- topological information, such as average length, average width, possible variations;
- geometrical information in the form of 2D shapes and sets of 3D points;
- where the vessel begins;
- where it terminates;
- important features (i.e. landmarks for radiologists);

- branches or tributaries;
- variations expressed as a percentage of cases;
- image features, such as as expected intensity, best x-ray view for the vessel;
- relations such as, anterior, lateral, anterolateral, posterolateral, posterior, medial, proximal, dorsal, in-front, behind, left-of, right-of.

The geometrical information slot in the blood vessel frame is used for displaying the vascular tree graphically. Automatic visualisation of the symbolic model enables visual inspection of the results making communication with the physicians easier while the model is being developed (Zrimec *at al*, 1995).

2.2 Brain Structure Frames

The relationship between brain features and the vasculature is established by the feeding vessels since the vessels supply or drain a particular part of the brain. These are organised as a hierarchy with the brain at the root with branches into large structures such as the left and right hemispheres. The brain is divided into two hemispheres. Each hemisphere branches into four main lobes continuing down to individual brain features. Each lobe is further subdivided. This part of the model is constructed from a co-planar stereotactic atlas of the human brain which is commonly used by radiologists and surgeons (Talairach & Tournoux, 1988).

Each brain feature contains *supply* and *drain* slots which indicate the blood vessels which supply and drain that feature. As we mentioned earlier, there are corresponding slots in the frames representing those blood vessels so that it is easy to determine which features they supply or drain.

2.3 View Frames

A symbolic representation of the slices, as depicted in the Talairach atlas, stores knowledge about spatial organisation. This information is similar to the information captured in MRI slices. Different sets of slices represent views from three different angles. Slices in each set are grouped according to their approximate location, e.g. ‘front or back’, ‘left or right’, etc. This segregation assists in matching slices with the structural representation of the brain.

The links between view frames and other kinds of frames indicate on which slices we would expect to find a particular brain feature or vessel and which vessel is expected in a particular slice. The frames also indicate which features surround each vessel.

2.4 Using the Model

We now return to the original problem of using knowledge to guide the interpretation of an x-ray image. Obviously, some kind of feature extraction is required to obtain information about the contents of the image.

The feature extraction system, currently under development, uses low-level image processing techniques to trace blood vessels. This generates line segments, organised as a tree structure, where each line segment corresponds to a blood vessel segment.

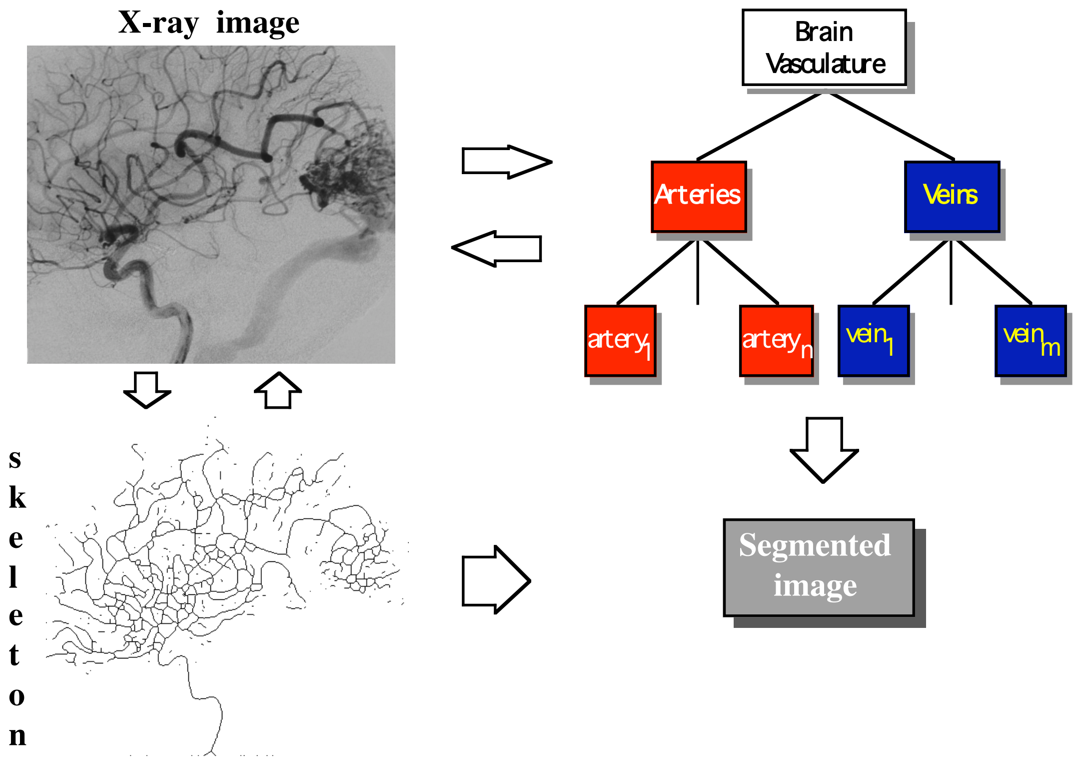

Figure 3. Using the model

The order in which vessel segments are extracted is governed by the frame hierarchy. The thick vessels are extracted and labelled before proceeding to the finer vessels. As feature extraction progresses, instance frames, representing the vasculature of the patient, are generated and compared with the frames of the model in order to recognise the object. The model can also be used to correct errors in the low-level processing. The tracing algorithm is far from perfect and may miss some blood vessels or add spurious line segments. The reference model of the brain can be used to “fill in” the missing anatomy of the patient’s model and filter out noise.

3 Learning and Refining Object Classes

In order to use a model for image processing, we have to have a model to begin with. Previously, we saw that a substantial part of the model can be derived from text books. However, these only describe idealised models of the vasculature or provide a catalogue of known variations. Books provide insufficient information to be able to recognise the many variations that a radiologist will encounter in practice. Therefore, one possible solution to building an adequate model of the cerebral vasculature is through some form of automatic or semi-automatic learning.

Although thousands of x-ray angiograms are taken each year, they are not routinely labelled and stored in a database for use by a machine learning program. Therefore two possible solutions to acquiring a refined model suggest themselves. The first is that an expert provides a small set of representative examples that have been carefully labelled.

The second approach is to allow the image processing system to attempt to recognise objects in the image and to be corrected by an expert when mistakes are made.

3.1 Learning

The task we address here is to label the blood vessels in the x-ray image. To build a model through machine learning, we first use the low-level image processing routines to produce a skeleton of the vasculature. The tracing algorithm then uses the skeleton to navigate through the grey-scale image and extract blood vessel segments. Recognition is normally done iteratively, first thresholding the image aggressively so that only the largest blood vessels are visible. This results in a relatively simple skeleton. Once the main blood vessels have been recognised, the image is thresholded again to allow slight smaller vessels to become visible. This is repeated until all the vessels have been labelled. This iteration is necessary since the skeleton of the complete image is very complex and also very noisy. Progressive thresholding allows us to contain the complexity and this also make it easier to eliminate noise.

A set of segments can be represented by a graph structure in which edges represent segments of a blood vessel and the nodes represent branching and bifurcation points and end points. Edges have associated information about the average diameter and average intensity of the vessel segment. The orientation of segments are stored as Freeman codes (i.e north, north-east, east, south-east, etc). An added difficulty is that since x-rays are two dimensional projections of three dimensional structures, some cross-overs and branches will appear as artifacts of the projection.

Although the surface representation of the knowledge in our system is in the form of frames, the frames are stored as predicates in Prolog's database. For many tasks, including learning, it is more convenient to operate directly on the predicates. The frame representation is more suitable for human readability. In the remainder of this section, we will only consider the Prolog representation. The segment graph is represented in Prolog's database by a set of predicates:

```
segment(S, Vessel, Directions, AverageDiameter, AverageIntensity, [Successors]).
```

where *S* is a unique identifier for the segment and *Vessel* is the name of the blood vessel to which the segment belongs. *Successors* is a list of segment identifiers for segments that are connected to this segment.

The segment predicates described above are automatically generated by the tracing algorithm. The expert must provide the following predicate:

```
blood_vessel(Vessel, InitialSegment, VesselType)
```

where *Vessel* uniquely identifies the blood vessel, *InitialSegment* is the identifier for the first segment in the graph and …*VesselType* is the class label for the concept to be learned.

Given these data, a relational learning system is able to learn a generalised description of a blood vessel from examples. *iProlog* incorporates a number of machine learning algorithms, including a relational learning algorithm that extends Plotkin's Relative Least General Generalisation (Plotkin, 1971). Since the output of the learning algorithm is a

Prolog program, it can be used as a recogniser for new blood vessel data. The program can also be transformed into a frame which is more readable to experts.

We now give an example of learning a simple labelling program for the internal carotid artery. Two typical training examples follow:

```
blood_vessel(v1, 1, 'ICA').
segment(1, v1, [n], 12, 22, [2, 4]).
segment(2, v1, [n, w, n], 11, 19, [3, 5]).
segment(3, v1, [ne], 10, 18, [6]).

blood_vessel(v2, 10, 'ICA').
segment(10, v2, [n], 15, 50, [20, 40]).
segment(20, v2, [w, n], 13, 25, [30, 50]).
segment(30, v2, [nw, ne, e], 9, 15, [60]).
```

The RLGG algorithm also requires negative examples to prevent over-generalisation. We omit the full training set and present a clause, to be constructed by the algorithm, that describes an internal carotid artery and its segments.

```
blood_vessel(V, S0, 'ICA) :-
    segment(S0, V, [n], Diameter1, Intensity1, [S1, S4]),
    segment(S1, V, Dirn2, Diameter2, Intensity2, [S2, S5]),
    segment(S2, V, Dirn3, Diameter3, Intensity3, [S6]),
    Diameter1 > Diameter2, Diameter2 > Diameter3,
    Intensity1 > Intensity2, Intensity2 > Intensity3,
    append(_, [w, n], Dirn2),
    member(ne, Dirn3).
```

This describes a blood vessel consisting of three segments in which the diameter and intensity of the segments consistently decrease (i.e. the blood vessel narrows and becomes less bright). The blood vessel always begins pointing north. The second segment must end with a turn from west to north and the third segment must contain a part that points north east. Note that the last condition should state that the third segment ends pointing either east or north east, but more examples will needed to discover this.

3.2 Refining the Model

The problem with learning from examples is that the set of examples must be sufficient to provide all the variations that the recogniser is likely to encounter. This is often difficult. A complimentary method for refining a model which is partially correct is to repair it when errors are detected. Compton (Compton & Jansen, 1988) has been successful in applying *ripple-down rules* to building knowledge bases incrementally under the supervision of an expert. In the PIERS system (Srinivasan et al, 1991), an endocrinologist supervised the performance of a program providing interpretations of laboratory assays. Radiology provides a similar environment in which an expert is available to critique the output of an expert system. It seems that the RDR methodology would be appropriate for this task. In the remainder of this section we discuss the possibilities for RDRs.

Since we already have a large body of text book knowledge that has been encoded into the system, the task for learning can be set up to refine existing knowledge rather than learning from scratch. RDRs were designed for refining knowledge.

The most significant obstacle for the use of RDRs is that, to date, they have only dealt with propositional data, i.e. the examples do not contain relational information as required in image understanding. This need not exclude RDRs from this domain. The main operation required by RDRs is that cases can be compared and differences isolated so that they can be used in exception rules. This is possible in first order representations. In addition, it must be possible to variablise expressions in the rules. Plotkin's Least General Generalisation may provide the mechanism for doing this.

Interaction with experts should be in the same medium that they normally use to process data. In the case of radiologists, that medium is visual, i.e. the x-rays. The user interface for an RDR system in this domain must be capable of displaying x-rays and allowing the radiologist to point to errors and attach comments to the display.

3.3 Current Status

The frame-based system described in sections 2 and 3 has been implemented although there is still considerable work to be done in improving low-level feature extraction. Work on learning is still in progress. The skeleton tracing has recently been incorporated into *iProlog* and coupling the output to the learning algorithm is proceeding.

iProlog currently implements a propositional version of ripple-down rules. These rules can be inserted into the demons of a frame. However, we are yet to adapt RDRs to labelling the blood vessel segments.

4 Conclusion

Medical image understanding requires diverse sources of knowledge to interpret visual scenes. The knowledge is represented in the form of object models that generate predictions during image analysis. The initial knowledge is too general and has to be augmented with additional domain specific knowledge.

We use anatomical, as well as other knowledge, as background knowledge and we acquire domain specific knowledge through learning. The model consists of symbolic, structural and geometrical information represented by frames. The frame representation provides a comprehensive description of the important cerebro-vascular forms and also provides a link between the image data extracted from the images and the symbolic knowledge. The frame representation is suitable for extending the existing structure and for storing additional specialised knowledge.

The internal representation of the frames is as predicates in Prolog's database. Since the hierarchical structure of frames maps well onto the hierarchical structure of blood vessels, frames have been found to be a natural medium for radiologists to express their expertise. However, for learning, Horn clause logic is a much more convenient representation. Thus, we use frames as a surface syntax for the user while performing learning on the internal, logical representation.

We use a relational learning system which, by generalising examples of labelled x-ray angiograms builds a labelling program. We wish to pursue the use of ripple-down rules which, by interaction with radiologists, will acquire specialised object descriptions. We believe that the approach being followed here will enable a medical image understanding system to gain experience in a specific domain and to improve object recognition with time, just as experts do.

Acknowledgements

Thanks to Geoff Parker for his valuable medical assistance in constructing the model. This work was partially supported by the Slovenian Ministry of Science and Technology.

References

Compton, P. J., & Jansen, B. (1988). Knowledge in context: A strategy for expert system maintenance. In *Proceedings of the Australian Artificial intelligence Conference*,

Goldstein, I.P., Roberts, B. (1977). Using Frames in Scheduling. In *Proc. Fifth Int. Joint Conf. on Artificial Intelligence*, August, 257-284.

Horn, W. (1992). Utilizing Detailed Anatomical Knowledgefor Hypothesis Formation and Hypothesis Testing in Rheumatological Decision Support. In *Deep Models for Medical Knowledge Engineering*. E. Keravnou, Elsevier Science Publishers B.V, 27-50.

Minsky, M. (1975). A framework for representing knowledge. In Winston, P (ed). *The Psychology of Computer Vision*, McGraw-Hill, New York, 211-277.

Nolte, J. (1993). *The Human Brain an Introduction to its Functional Anatomy*, Mosby-Year Book, Inc., Missouri.

Plotkin, G. D. (1971). A further note on inductive generalization. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence 6*. New York: Elsevier.

Sammut, C.A. (1997). *iProlog*. ftp://ftp.cse.unsw.edu.au/~claude/src/iprolog.tar.gz.

Srinivasan, A., Compton, P., Malor, R., Edwards, G., Sammut, C., & Lazarus, L. (1991). Knowledge Acquisition in Context for a Complex Domain. In *Proceedings of the European Knowledge Acquisition Workshop*, Aberdeen:

Talairach, J., Tournoux, P. (1988). *Co-Planar Stereotaxic Atlas of the Human Brain, 3-Dimensional Proportional System: An Approach to Cerebral Imaging*., Time Medical Publishers, Inc. Georg Thieme Verlag.

Zrimec, T., Mander, T., Lambert, T., Parker, G. (1995). 3-D Visualisation of the Human Cerebral Vasculature, *SPIE 95 Medical Imaging: Image Display*, Vol 2431.