

The Use of a Computerized Brain Atlas to Support Knowledge-based Training in Radiology

Serge Garlatti^a, Mike Sharples^{b*}

^aDépartement Intelligence Artificielle et Sciences Cognitives, Ecole Nationale Supérieure des Télécommunications de Bretagne, Technopôle de Brest Iroise, BP 832 29285 Brest, Cedex, France.

^bSchool of Electronic and Electrical Engineering, The University of Birmingham, Edgbaston, Birmingham B15 2TT, UK.

*Corresponding author. Tel.: +44 121 414 3966

Fax: +44 121 414 4291 Email: m.sharples@bham.ac.uk

Correspondence address: School of Electronic and Electrical Engineering, The University of Birmingham, Edgbaston, Birmingham B15 2TT, UK.

Abstract

Trainers of radiologists face the particular challenges of teaching normal and abnormal appearance for a variety of imaging modalities, providing access to a large appropriately-indexed case library, and teaching a consistent approach to the reporting of cases. The computer has the potential to address these issues, to supplement conventional teaching of radiology by providing case-based tutoring and diagnostic support based on a large library of images of normal and abnormal anatomy, described in a consistent terminology. The paper presents a new approach to computer-based training in radiology that combines a knowledge-based tutor with an on-line medical atlas. It describes two existing computer systems, The MR Tutor and ATLAS, and discusses the medical, computational, epistemic, and pedagogic issues involved in developing a combined Atlas-Tutor. Integrating an atlas with a training system could significantly improve the teaching and support offered, but practical difficulties include the need to merge knowledge representations and to incorporate techniques for registering atlas plates on images that exhibit abnormalities. The paper addresses these problems, and concludes by indicating how the Atlas-Tutor might be employed in practical radiology training.

Key words

Knowledge-based training, computer-based atlas, knowledge representation, magnetic resonance imaging, neuroradiology

1. Introduction

A recent study of medical training in the United Kingdom [1] identified many problems with current teaching methods and proposed a curriculum that would include: a greater emphasis on developing skills rather than communicating facts; introduction of self-directed and problem-based learning; decentralised teaching; fair and objective assessment. Trainers of radiologists face not only this general need to reform the medical curriculum, but also the particular challenges of teaching normal and abnormal appearance for a variety of imaging modalities, providing access to a large indexed case library, and teaching a consistent approach to the reporting of cases.

Two separate teams, in the UK and in France, are addressing these issues by developing, respectively, a computer-based training system for MRI neuroradiology [2, 3] and an on-line atlas of the brain [4]. These systems have been developed to prototype stage and are currently being evaluated for medical use. A general aim of both these research groups is to provide an environment that supports active learning and enquiry, that can adapt to the user's needs and abilities and can cover a range of radiological tasks including image interpretation, diagnosis, and surgery planning.

This paper discusses the benefits to be gained by merging the two approaches, to provide a computerised brain atlas in support of radiology training. The paper reviews existing training systems for radiology and presents the need for a case-based training and diagnostic support system. It describes the current MR Tutor system and concludes that while it is able to offer flexible training and support, it could be enhanced by the addition of an integrated atlas of the brain. The paper then surveys computer-based atlases and describes the Atlas project to develop a knowledge-based brain atlas as part of a medical decision support system. Integrating such an atlas with a training system could

significantly improve the teaching and support offered, but practical difficulties include the need to merge knowledge representations and to develop new techniques for registering atlas plates on images that exhibit abnormalities. The paper addresses these problems and discusses how the combined Atlas-Tutor might be deployed in radiology training and some practical problems of integrating the two technologies.

2. Training and the development of expertise in radiology

Lesgold and colleagues have observed and analysed cognitive processes in radiology [5]. They conclude that expert radiologists carry out a multi-stage process of interpretation. On first seeing a film an expert radiologist rapidly invokes a mental schema that covers the salient features of the image and by comparing this with schemata for normal anatomy, detects areas of abnormality. This results in one or more tentative diagnoses. The radiologist then 'tunes' the interpretation by searching for perceptual features that have been missed, by referring to clinical information on the patient, by compensating for technical defects in the film, and by making inferences about the cause and progression of the disease. Lastly, the findings are articulated as a verbal report.

The ability to integrate context-specific visual schemata with more general problem-solving and biomedical knowledge (knowledge of anatomy, aetiology, histology, biochemistry etc.) is central to the diagnostic process of radiology [6]. In a series of experiments on the role of biomedical knowledge in clinical reasoning, Boshuizen and Schmidt [7] have shown that expert clinicians, unlike novices and intermediates, do not apply biomedical knowledge as a distinct stage of reasoning. Rather, as a result of a long process of skill acquisition, they have learned to encapsulate it into their diagnostic process. The development of knowledge integration is an active process. The learner must

acquire and articulate the general biomedical knowledge, diagnose, discuss and reflect on individual cases as they are encountered, make connections between situated experience and general medical knowledge, and repeat this process over many cases until the biomedical knowledge becomes an integral part of diagnostic reasoning.

Training schemes for radiology can assist this process of knowledge integration by a) providing trainees with exposure to a large well-structured archive of appropriately indexed and described cases and b) helping trainees to integrate the experience they gain from reporting the cases with more general biomedical knowledge.

The ability of training institutions to build archives of teaching material is hindered by the fact that there is no agreed terminology for describing radiographic images. Standard terminologies are widely used for disease categories and anatomical structures, but there is no similar language to describe abnormal appearance. Each radiologist uses different descriptive terms, or worse, similar terms with different meanings. Although there is no standard method for describing abnormalities, there is widespread agreement on the need to develop a more structured approach to reporting, so that radiologists can learn a consistent language with terms that have been precisely defined. There is also a need to annotate the case images with precise anatomical borders so that trainees can acquire perceptual schemas that delineate correct anatomy.

Both extensive individual tuition and the provision of a large consistently-described archive are beyond the resources of conventional radiology training, but they could be provided by computer-based training.

3. Computer-based training in radiology

Computer-based training in radiology offers a significant enhancement to the current methods of lecture, tutorial, apprenticeship and self-study, by allowing trainees to browse and search through libraries of case images, by offering problem-based tutorials, and by giving trainees rapid access to extensive textual and visual information. Some recent systems for training radiologists are powerful and sophisticated, providing 3-D visualisation of anatomy [8-11], hypermedia reference aids [12], case-based instruction [13] and training via the Internet [10, 14-16]. However, almost all the current training systems are reactive. They respond to answers, queries or probings, but they have only a very limited capability to adapt to the specific needs or skills of the user, to cooperate in performing a diagnosis, or to engage in a tutorial dialogue.

We can describe computer-based training and work support systems in terms of a spectrum from passive and reactive to active and adaptive. The main types of system along this spectrum are:

- a) reference aids,
- b) computer-assisted instruction,
- c) simulations,
- d) case-based tutoring and diagnostic support.

It is easiest to implement systems for radiology training and support at the reactive end of the scale, by means of computer-assisted instruction, or through reference aids with no direct teaching. These systems do not require a detailed computer-interpretable model of the user or the teaching material. The more difficult, but potentially more fruitful, problem is to design a system that adapts to users' needs and abilities and engages them in problem-solving activities to test and diagnose their process of radiological interpretation.

3.1 Reference aids

Reference aids have been developed for a number of imaging modalities including CT and MR [17]. Typically, they allow the student to browse through a library of images organised according to pathology or position in the body, and to call up textual annotations. Whereas the simpler systems are little more than image libraries annotated with informative labels, the more advanced ones offer extensive control over the structure and presentation of the images, including the ability to present 3D reconstructions from MR slices, along with structured training texts. They offer a valuable backup to human or textbook teaching, particularly if the teaching is linked explicitly to the specific reference material, but they are limited by the shallowness of relation between the images and the textual material. The descriptive labels and pre-prepared texts do not form a systematic representation of knowledge that could be interpreted by the computer to provide active teaching, or to answer complex queries.

3.2 Computer Assisted Instruction

In computer assisted instruction (CAI) the knowledge and expertise is pre-compiled into 'frames' of teaching information. Each frame typically consists of a short piece of teaching material and a multiple-choice question. When it is presented on the computer screen, the trainee reads the teaching material and selects an answer. If the answer is correct the system offers confirmation; if it is wrong, the system provides supplementary teaching and a further set of questions. Within this general format CAI can be sophisticated in structure and appearance. It can provide multiple pathways between the frames and training material that draws on illustrations, images and animation. An example is MITS, the Medical Image Teaching System [18]. A typical MITS teaching frame presents one or more images, a question relating to identifying abnormalities, and a list of possible answers.

A CAI system is limited by its inability to know what it is teaching, to a level that would allow it to answer questions from the learner or to support learner-directed study. For example, unless the responses have been anticipated in advance, it cannot explain why it has designated one answer as correct and another as wrong, nor can it track a student's developing expertise, to respond in a style and detail that matches the learner's current understanding.

3.3 Simulations

Simulations overcome the pedagogic limitations of CAI, by providing an exploratory environment rather than a lesson. This may range from an emulation in software of a piece of medical equipment such as an MR scanner, to a 'virtual patient' on which a student can perform tests or surgery. An example of the former is SimBioSys [19], a dynamic simulation of an intensive care unit, with a bedside monitor displaying ECG waveforms, pressure waveforms and other clinical data. The student observes, diagnoses, and treats the monitored patient using a stock of clinical tools, such as a ventilator and infusor, provided by the system. As an example of the latter, Marshall University School of Medicine has developed The Interactive Patient [20] that allows a medical student to simulate an encounter with a patient, carrying out a physical examination, requesting additional history and reviewing laboratory data and X-rays.

Simulations can be powerful aids to learning, offering detail and realism without danger, under the control of the learner, and providing opportunities for incidental learning. A simulation alone, though, does not provide teaching or direct cognitive support; that must be given either by a human teacher or a training system.

3.4 Case-based training and diagnostic support

To provide training of a quality that approaches human experts, a system must overcome the limited responses of CAI and the shallow knowledge of reference aids, by storing structured representations of knowledge about the domain and the student which it calls on to generate sequences of teaching actions and to diagnose learner misconceptions. The teaching actions that a human teacher of radiology carries out can broadly be categorised as:

instructive: e.g. teaching the meaning of a term, or giving an introductory lesson;

indicative: e.g. pointing out a relevant case, or highlighting anatomical structures on a case image;

remedial: correcting a misconception;

interrogatory: setting questions to test the student's knowledge;

managerial: e.g. presenting a new case;

supportive: assisting browsing, or offering encouragement.

A knowledge-based tutor could, in theory, offer all these types of teaching through a combination of structured representations of biomedical and situated knowledge, rule-based teaching strategies, and a dynamic model of the student's understanding. In practice, the difficulties of formalising medical knowledge and designing generative teaching for a domain that combines multiple media mean that few such systems have yet been developed for medical imaging.

Despite the difficulties of design and implementation, the combination of case-based training and knowledge-based diagnostic support offers the possibility of a powerful and flexible training aid. A team comprising software

designers, medical statisticians and neuroradiologists is developing such a system as part of the MEDIATE project.

3.5 The MEDIATE project

MEDIATE is a joint project between the University of Sussex, De Montfort University, Leicester, the University of Birmingham, and the Institute of Neurology, London. Its aim is to address a need identified by the radiology specialty for a more structured approach to reporting and training. One outcome of the project is the MR Tutor, a training system to assist radiologists in interpreting MR images of the brain, particularly images presenting diseases that are acknowledged to be difficult to differentiate.

The MR Tutor combines an active tutor with a responsive support system. It is based on a constructivist approach to learning, in which the trainee is helped to acquire a well-structured approach to describing abnormal features of images by engaging in an active process of case-based reporting. The MEDIATE team has devised a structured image description language (IDL) for MR neuro-images suitable for a wide range of imaging sequences. The language is image-based, covering the visual appearance and anatomical locations of the image abnormalities. An archive of around 1300 cases, fully annotated using terms of the IDL and accompanied by clinical information and confirmed diagnoses, forms the knowledge base of the Tutor.

A trainee can browse through the cases and either the trainee or the system can call on images for teaching. The trainee identifies and describes the abnormal features in the case images using the IDL terminology and the system provides guidance towards a full and correct interpretation. A feature of the Tutor is its use of the statistical technique of multiple correspondence analysis [21] to generate an 'overview plot' that shows the distribution of cases by both appearance and position of abnormal features in the images for a given

disease. It provides a direct display of similarity (in terms of distance between displayed points) and typicality (distance of a point from the centre of the plot) and allows the user to access the image archive by interacting with points on the plot (see Fig. 1). By viewing overlaid plots for two or more diseases, the trainee can gain a general impression (a gestalt) of the distribution of cases across diseases and can identify particular borderline cases that may be difficult to diagnose.

FIGURE 1

The IDL has been developed over some ten years, involving knowledge elicitation from an expert neuroradiologist, to ensure that, as a language, it is precise, necessary and sufficient to support accurate diagnosis of confusable diseases [22]. The validity of the overview plot has been assessed [23] by asking radiologists to view a set of cases associated with a plot for one disease (gliomas) and then to place new cases in the appropriate position on the plot. The results show a good correlation between the placement of the same cases by experienced radiologists and by the multiple correspondence analysis algorithm.

A heuristic evaluation of the MR Tutor was carried out to assess its usability and its general effectiveness as a training system. The evaluation involved two groups of subjects, one of people with considerable experience of system design and Human-Computer Interaction, and the other of experienced neuroradiologists. On a 5 point scale of overall satisfaction, the mean rating of the HCI group was 3.96 and that of the radiologists was 4.13. In general, both groups reported that the system was stimulating and easy to use. The HCI subjects reported that the system conforms to good practice in interface design

and many of the radiologists indicated that it could provide a useful adjunct to their teaching. Both groups suggested minor improvements to the interface and functionality.

One general limitation of the MR Tutor is its inadequate access to structured knowledge and to illustrations of normal and abnormal anatomy. For example, if the trainee fails to indicate one or more structures affected by a lesion this may be because either:

- a) the trainee has misinterpreted the extent of the lesion (perhaps because of its similarity in appearance to normal tissue), or
- b) the trainee is aware of the extent of the lesion but cannot distinguish or name the anatomical structures that it affects.

One way to distinguish between these two types of misconception would be for the Tutor to ask the student to mark the extent of the lesion. If the student were unable accurately to outline the lesion then the tutoring could concentrate on lesion identification. Alternatively, if the student were able to outline the correct border of the lesion then the system could infer that the trainee lacks knowledge of normal anatomy, and that it should provide remedial instruction about the affected anatomy. To do this, the system would need to call on the names and regions of the anatomy affected by the lesion and be able to display these in the correct location on the case image. To deal with further misconceptions the system might also need to have knowledge of the adjacent anatomical structures, the sizes, shapes and radiological appearance of the affected anatomy, and their appearance in relation to the lesion. The current MR Tutor is unable to carry out this quality of teaching.

More generally, to provide an informative response, the system may need to display the borders of anatomical regions in register with the case image, for different image sequences and orientations, and to provide information about

the histology, biochemistry and functional properties of normal and abnormal anatomy. A computer-based atlas might be able to offer such assistance. The integration of an atlas with a case-based radiology training system offers the possibility of enhanced quality and flexibility of training.

4. Computer-based atlases of the brain

Neurosurgeons, neurologists, and neuroradiologists consult atlases of the brain to supplement their knowledge of anatomy and brain functioning in support of medical decision-making. Studies of radiological interpretation [5] suggest that expert radiologists hold their anatomical knowledge as 3-D schemata which they can invoke to visualise 2-D radiological appearance, and most atlases provide a combination of annotated 2-D images and illustrations of the 3-D brain [24-27]. Computerised atlases are still not widely used by the medical profession, which sees them as inconvenient and lacking authority, but in recent years they have evolved from basic computerised maps of the brain into sophisticated medical decision support systems. In this section we indicate the scope and purpose of brain atlases and we describe the new facilities that a computerised brain atlas can provide.

4.1 Scope and purpose

The main purpose of a brain atlas is to support the integration of case image and general anatomical knowledge. Atlases can assist with identifying anatomical structures in case images, establishing relationships between different brains by comparing their anatomy, and specifying the functional properties of visible structures. In radiology practice atlases are used for:

interpreting images: to assist in visualising or outlining anatomical structures in 2D or 3D on case images, and in naming the anatomy;

diagnosis: for example to identify the anatomy responsible for generating epilepsy by correlating brain function with electrographic and clinical findings;

surgery planning: to identify the names, extent and function of relevant anatomical structures, in order to find a safe path for surgical intervention and to predict the functional consequences of surgery.

Printed radiological brain atlases normally indicate the major regions and some more obvious smaller structures. Radiologists have been slow to employ the more sophisticated and detailed knowledge of tracts, nuclei and their functions that is familiar to a good neurologist. A computer-based atlas may be able to provide some of this invisible detail and functional information overlaid on radiological images.

Most current computer atlases are reactive systems that respond to a user's requests by displaying annotated illustrations and biomedical information. An alternative is for the atlas to form the core of a more active environment, that supports user's needs and abilities and augments their skills.

In the following sections we shall outline two types of reactive atlas, brain maps and decision support systems, and describe the advantages of a computer implementation over a printed version. This provides a foundation for discussing the design of more active atlas environments for radiology.

4.2 Computer-based maps of the brain

The simplest type of computer-based atlas is just a map of the brain containing illustrations scanned or copied from printed atlases or medical books along with a terminology, a coordinate system, reference symbols, and descriptors to indicate salient features [28]. It presents information and offers the ability to view and compare brains in pictorial and symbolic form.

More sophisticated brain maps offer both 2-D and 3-D views, visualisation tools such as cutaway illustrations, and the ability to compare an area of the brain from different perspectives such as neuroanatomy, cerebral function and blood flow. The atlas can also contain images captured by CT or MRI, and several images in register can be merged to form a 'average' brain.

The text and graphic representations can be combined and displayed through a multimodal interface, with active links from areas in the image to descriptive text. By mixing symbols with images, a computer atlas can enhance and facilitate comprehension. For instance, a textual description of the relationship between two gyri can be matched to a display of a 3D volume showing the two structures outlined and named. Techniques for hypermedia referencing and automated indexing allow the atlas data to be accessed by multiple routes.

To assist visualization, programs have been developed [29] that build sets of case slices into 3-D volumes. Once brain data is stored as a digitised volume, the computer can then display slices from any viewpoint, or can display the brain as a 3-D object that can be rotated and cut away to show its internal structure. A dynamic annotated 3-D view of a case offers clear advantages over a paper atlas with its limited set of views, particularly for surgery planning. Other visualisation tools can alter the parameters of the image such as brightness and contrast, omit detail and highlight features.

Such computerised maps of the brain are the equivalent of reference aids in computer-based training. They can be used for the same teaching purposes and suffer the same limitations. In terms of the process of knowledge articulation, they can help a radiologist to articulate a global framework of biomedical and anatomical knowledge, but are less useful for reflecting on individual cases or for making connections between situated knowledge and the

global framework. They lack dynamic tools to relate case images to normal anatomy or to interpret inter-individual variability, for example by registering case images over reference illustrations and thus outlining and labelling the anatomy from any viewpoint. This task has to be achieved by mentally comparing a case image and the corresponding atlas plate.

4.3 Decision Support Systems

Given the limitations of map atlases, an alternative is to develop a set of tools based around an atlas that can provide support in medical decision-making. However, the types of decision-making that require use of a medical atlas cannot be formalised; they are what Simon calls “unstructured decisions” [30]. Radiologists rely on heuristics rather than well-specified algorithms, and can backtrack if their strategy seems unproductive. Decision support systems are designed for just such situations. The early ones were developed as expert systems, in the expectation that it would be possible to produce a system that performed at least as well as a human decision-maker. More recent ones recognise the difficulty of formalising medical knowledge, and offer a range of tools to support a human expert in reflecting on situated experiences, externalising the steps in solving a problem, and integrating perceptual, causal and anatomical knowledge [6, 31, 32]. Taylor [33] gives a review of computer aids for decision-making in diagnostic radiology. It covers image databases, numeric and statistical systems, expert systems, image processing systems and image understanding systems. Although the review states that conventional atlases may not provide an adequate basis for a decision support tool, it does not indicate how they might be adapted for the purpose. A decision support system based on a brain atlas can provide more than just a set of illustrations, it can offer tools for data handling, image manipulation and plan formulation.

An image database linked to an atlas can store patient data, including 2-D and 3-D images, that can be searched or browsed by means of hypermedia links [4, 10, 14, 15, 34-37]. A database can also store statistical information about the variation in structure across a number of cases or reference brains. An image database designed for decision support is Cohn et al.'s AXON [38]. Each image in the database is associated with a set of keywords labelling the lesions and the disease. From a keyword and a condition the system will search through a hierarchy of disease types and display images matching the condition.

Another facility that a decision support system can offer is a set of tools for image annotation, enhancement and interpretation. These include drawing packages for outlining and naming regions on case images, filters such as histogram equalisation and Gaussian blur, and techniques such as a 'warping model' for matching atlas plates to case images. A warping model is an algorithm that employs rigid (linear) or elastic (non-linear) deformations to align appropriate slices from atlas illustrations in register with case images. If successful, the structures outlined and labelled on the reference illustration will be accurately superimposed on the corresponding ones from the case image. The same technique can be used to register images of low spatial resolution, such as PET, SPECT and MEG, with high resolution CT or MRI reference images to locate functional activities on the brain [39-45]. Image interpretation can also be carried out symbolically, using model-based image interpretation and symbolic inferencing techniques to identify segmented regions on a case image [46]. Both warping models and symbolic interpretation are limited in their reliability, particularly for peripheral regions of the brain where there is greater inter-individual variability. But new methods of elastic deformation [47] offer the possibility of fully-automatic registration, providing the case image is sufficiently similar to an annotated reference image.

Knowledge of diseases and lesions in the AXON system is stored as frames (the word 'frame' is used here in the AI sense, and not with the meaning used earlier in the section on computer-aided instruction) that represent properties for a given class along with pointers to its immediate superordinate class. Such a structured representation of knowledge enables properties from the more general frames to be inherited by the more specialised ones. For example, the tuberculosis frame inherits properties from the mycobacterial infection frame which itself inherits properties from infectious diseases. The lowest level of the hierarchy represents the individual cases known to the system. This approach can be generalised to store knowledge in support of diagnosis for topics such as neuroanatomy and neurochemistry. Anatomical structures may be described by properties such as location, shape, size, grey-level and function. These attributes can correspond to a single reference brain, or to a statistical or probabilistic amalgam of multiple brains. The knowledge base can also store explicit information on the topology, orientation and spatial relations between structures. Further causal and domain knowledge from disciplines can be added to assist diagnosis [48].

With the ability to manipulate data in digital form, computerised atlases offer new opportunities to support decision making in radiology. In addition, they have advantages over printed atlases in enabling a user to modify and extend the information, and to access information at a distance. An authoring tool allows users to extend the atlas by adding new plates [35, 49], new images as reference cases to illustrate the variability of normal anatomy [50], or new data and knowledge sources [11, 48, 51][10, 45, 48, 49]. Information on a computer atlas can be distributed across a number of sites and accessed at a distance via a wide area network.

The use of atlases as brain maps and as decision support systems offers the possibility of making background knowledge available to radiologists in situ, which could form part of an integrated radiologist's workbench [52]. But they are reactive systems that can only respond to user's requests. A more fruitful approach may be to design an active environment that augments their skills and adjusts itself to each user's needs and abilities.

4.4 The Atlas project

Atlas is a joint project between IASC Laboratory at Télécom Bretagne in Brest and the Laboratoire SIM from the University of Rennes I to design an active atlas-based decision support system for neurosurgeons and neurologists in treating human brain pathologies. It takes a mixed-initiative approach to human-computer interaction, where the user can interrogate the system and the system can offer hints to the user. The aim is for the control and the tasks to be distributed between user and system in an adaptive manner so that the computer acts as a partner in problem solving. To provide an active environment, a decision support system also needs to be supplemented by a user model that it can access to generate intelligent help and context-specific advice.

Atlas provides an environment to support image interpretation, and future versions will also offer support for surgery planning, diagnosis and learning about the brain. The computerised atlas consists of a brain map, a database, imaging tools including procedures for performing 2D and 3D segmentation, 2D and 3D displays, linear and non-linear warping models for registration of the atlas on case images, and a knowledge base system as a decision model to support medical tasks. These components (computerised map, imaging tools, knowledge base and data base) are implemented in different computer environments (C++ libraries, Y3 — an Object oriented frame-based language

[53], and O² — a commercial object oriented data base [4, 34, 54, 55]) and will be linked together through a hypermedia interface which supports information retrieval and will manage the cooperation between human and computer (Fig. 2).

FIGURE 2

The brain map consists of a set of 2D digitised and visualizable slices, based on the Talairach-Tournoux printed atlas [24], along with texts, drawings and diagrams from printed atlases and medical books, and 3D volumes from different modalities derived from volunteers and patients. An object-oriented database supports the manipulation of complex objects and schema evolution which includes removing, merging and splitting of object classes. An instance migration function has been implemented that reorganizes the distribution of the instances between the classes. The project has developed a new semantics of schema evolution operations that extends the functions of current object-oriented databases with the specific needs of atlas applications.

The imaging tools provided by Atlas can perform segmentation tasks (such as brain segmentation, grey/white matter labelling, and sulci modelling) [44, 45], display 3D images by using volume and surface rendering tools [29], and perform linear and non-linear registration [56]. The project has investigated different approaches to solving the difficult problem of finding correspondences between brains of different subjects. Usually this task involves the computation of a generic warping model that can assign a 3D deformation vector to every voxel labelled as brain matter. A deformation vector enables the program to match a voxel of the base (atlas) volume to one and only one voxel of the matched volume. This global minimisation procedure assumes that the brain

behaves as an elastic object, and that the brain deformation can be modelled by prior knowledge. This prior knowledge can be expressed through mechanical models [57] or by statistical methods (in a Bayesian framework for instance) [57-59]. The deformation can also be controlled by matching up landmark structures on the base and matched volumes. Following this route, the Atlas team has introduced anatomical constraints on the non-linear procedures by using grey/white matter and some cortical sulci as reference points [60].

A major design aim is to provide a knowledge base that integrates multiple knowledge sources. The knowledge is organised by type (such as knowledge about tasks, methods and domains) and by level of abstraction (task-level knowledge and more general knowledge about disciplines) into loosely coupled parts with a unified representation. The partitioning of domain knowledge is crucial because it is not possible to design a single ontology that includes every aspect required to model the world. Dividing the world into distinct knowledge sources makes the knowledge base easier to reuse, understand and update. Different parts of the knowledge base can then be used for different purposes requiring domain knowledge, for instance, assistance with problem solving, intelligent help and case-based training.

The user interacts with the system via a hypermedia tool that offers flexibility of access and cooperative support. It manages the information search and retrieval (for cases, slices, texts, drawings, knowledge) and calls up the computation tools such as the warping algorithms, visualisation algorithms and problem solving methods. To support computation, hypermedia has to be dynamic [61, 62]. To take into account user's needs and abilities, a context-sensitive hypermedia system has been designed as the first step towards a cooperative environment [63, 64]. At present, it only supports data base queries and information retrieval for navigation (see Fig. 3), using stored knowledge of

domains, atlas plates, texts and drawings. In a future version, image tools (such as warping models and 3D visualization of patient data) will be available within a cooperative environment based on the dynamic and context sensitive features of hypermedia.

FIGURE 3

Atlas provides a supportive environment for expert neurosurgeons and radiologists, but it is not a training system. It does not possess a model of the student's skills and misunderstandings, nor an explicit and adaptable teaching strategy. But its modular design, based on a set of knowledge-based tools to support image interpretation, provide the basis for extending it from an expert decision aid to a combined training and support environment.

5. The design of an Atlas-Tutor

The previous sections have identified the need for a training system in radiology to be supplemented by an atlas, and they have indicated how an atlas can be incorporated into a decision-support system. The remainder of this paper discusses the ways in which a knowledge-based training system for radiology can be combined with a computer atlas.

Merging the software of two major projects presents many problems of interface design, knowledge reconciliation, and mutual interaction between the two systems. The easiest approach is to provide an atlas as a separate software package, for the user to call up as a reference aid to assist in identifying and describing normal anatomy. But the advantage of a tightly-coupled tutor and atlas is that both the trainee and the system can call on the

same resource, with the computer tutor accessing symbolic and pictorial knowledge of normal anatomy to support its teaching actions.

5.1 Teaching issues

In section 3.4 we described six types of teaching action that a radiology teacher might employ: instructive, indicative, remedial, interrogatory, managerial, supportive. A computer atlas can support each of these types of teaching. A human teacher may find it too unwieldy to call up a computer atlas during a lecture or tutorial, but a computer-based tutor can invoke the appropriate atlas feature on demand, while maintaining the context of the teaching. Having a range of tactics and resources to call on allows a computer-based tutor to make dynamic adjustments to its teaching approach, to suit the ability and learning needs of the individual trainee.

5.1.1 Instructive teaching

A training system can employ an atlas to teach basic anatomy. For instance, it can give refresher lessons, asking the student to point out named parts on the atlas or to name a given region. This is the easiest teaching action to implement, since it just adds an additional computer-aided instruction module to the case-based tutor. It only requires a computerised map indexed by structure and domain knowledge that could be interpreted by the computer, parts of which can be linked to teaching information and accessed by the teaching component. Although the implementation is straightforward, difficulties arise in providing a comprehensive and accurate database of anatomical knowledge. For example, some named regions and structures in the brain can only be identified by reference to adjacent or more distant structures, and some have no identifiable boundaries.

5.1.2 Indicative teaching

If the system's student model detects that the trainee is making general errors in recognising regions of an image, or if the trainee requests information about its visible structures, then the system can respond by annotating case images with anatomical structures taken from an atlas. The case images could be pre-prepared with annotations, but for a large case library marking up and labelling the structures on each image would be laborious. A better method may be to apply a warping model that displays an appropriate plate from an atlas in register with the case image. If an atlas plate can be registered with a case image, the image can then be annotated as needed with atlas illustrations, or the system can label any region of the case image with its anatomy derived from the atlas. Although successful warping algorithms have been demonstrated for normal anatomy, they may not be effective for cases displaying lesions.

5.1.3 Remedial teaching

Central to radiology is the ability to identify and name the areas affected by disease or other abnormality. There are four general types of conceptual/perceptual error that a trainee might make in describing the location of an abnormality on a case image:

Omission: the trainee fails to name an anatomical structure covered by the abnormality.

Addition: the trainee names a structure that cannot be detected as abnormal from the image (but can be detected on other images, or by other methods), or one that is not abnormal.

Translation: the trainee describes the abnormality in an incorrect position (i.e. a combination of addition and omission).

Congruence: the trainee is incorrect in judging whether an abnormality conforms to the borders of an anatomical structure.

In general, these errors may be caused by:

- a) the trainee failing to detect the abnormality or misunderstanding its extent or type (i.e. lacking diagnostic knowledge), or
- b) the trainee knowing the extent of the abnormality, but mistaking the anatomy it covers (i.e. lacking anatomical knowledge).

A linked atlas enables the system to distinguish between the two causes of student error. If it has access to an atlas registered with the case image, the tutor can request the trainee to draw round the boundary of the abnormality and can then call on the atlas to identify all the anatomical structures invaded or displaced by the marked region. The general algorithm is as follows:

1. Identify the correct borders of all the abnormalities (e.g. by having an expert mark up abnormalities on the case images, as part of the teaching resource).
2. Use a warping algorithm to register an appropriate reference image from the atlas over the current case image.
3. For each abnormality, determine which anatomical structures are affected, the extent of the coverage, and the type of change (invasion or displacement).
4. Generate a symbolic description of the relation of the abnormality to the anatomical structures, which can be matched against the trainee's description and used to guide the remediation.

By this means the system could distinguish between the two types of error, (a) and (b) above, and offer remedial teaching. If the border of the abnormality drawn by the trainee does not match that provided by the expert (to within some tolerance), then the trainee lacks diagnostic knowledge and the system can tutor around the issue of detecting and describing abnormalities. If the border drawn by the trainee is approximately correct, then the trainee lacks anatomical

knowledge and the tutor can offer remedial advice around the mismatch between the trainee's description of the anatomy involved and that detected by overlaying the atlas on the area around the abnormality. The tutor can also use the atlas to highlight the structures mentioned by the trainee, show how they are affected by the abnormality and give some explanations to the user.

Many issues need to be addressed in the design of such tutoring, including accurately registering an appropriate atlas plate on the case image, identifying the type and extent of the deformation to anatomy caused by the abnormality and, since many anatomical structures have no clear boundary, providing a response with appropriate tolerance.

5.1.4 Interrogatory teaching

There are a number of ways an atlas could be called on to test the trainee. For example, the system could:

- a) ask the trainee to name the anatomy affected by a lesion;
- b) indicate a region on a case image and ask the trainee to name the anatomy;
- c) request the trainee to draw in and name the boundaries of all the major anatomical structures visible on an image;
- d) select a region and then ask the trainee to indicate the (invisible) structures that are known to lie within it.

If the system is able accurately to overlay an annotated atlas plate on any given case image, then it can generate all these types of questions and their answers by

- a) retrieving a list of the names of the anatomical structures covered by the lesion, asking the student to name them, and matching its list against the structures suggested by the student;

- b) selecting a random region on a case image or one generated from knowledge of the anatomy, retrieving a list of the associated structures, asking the trainee to name them, and then matching these against the trainee's response;
- c) comparing the boundaries of the structures drawn by the student with those on the registered atlas plate, and either reporting (e.g. in terms of areas covered) or displaying the differences;
- d) retrieving information on the tracts and nuclei that are known to lie within a given region of the brain.

5.1.5 Managerial teaching

If the training system has access to a model of the student's knowledge (derived from earlier interrogatory and remedial teaching) in terms of anatomy that the student finds difficult to locate or label, then the program can select a case involving those particular anatomical structures. Alternatively, if the system has a more general model of anatomical structures that are generally hard to identify or are regularly confused, then it could call on standard techniques of concept tutoring [65] to generate an ordered sequence of cases involving the confusable structures.

5.1.6 Supportive teaching

The system can support student-directed learning by making an atlas easily available online, along with a variety of tools for viewing the atlas from any projection, visualising anatomy in 3-D, registering the atlas plates on any case image supplied by the student, and linking the anatomical structures to descriptions of function, for example of the optic or aural system.

5.1.7 New types of teaching

Adding an atlas not only increases the range of possible teaching responses, but also offers the possibility of new types of teaching. Two of these are described here: structured reporting, and hypothesis-driven remediation.

A structured approach to reporting, based on a systematic approach and language for describing images, enables radiologists to exchange their findings in a common terminology, with defined referents for each term [22]. The Image Description Language used for the MEDiate project provides sets of properly-defined terms appropriate to diagnostic reporting. As a medium for teaching it has the advantage that the trainee can learn a canonical set of terms supported by definitions and examples, through a well-structured curriculum. An atlas provides a means of extending the definitions of anatomical terms by relating them to regions on reference images. So, for example, if a trainee wants to know the usage of cerebral white matter all areas can be indicated on the atlas, along with a grey-level chart showing the variation in cerebral white matter intensity for a particular sequence.

Hypothesis-driven remediation is an effective method of tutoring and has been incorporated in successful computer-based training systems including the SOPHIE series of tutors for electronic troubleshooting [66]. The basis of hypothesis-driven remediation is that if a trainee offers an incorrect response, then the system tutors about the difference between the student's hypothesis and the correct answer. For example, if a trainee misidentifies an anatomical structure or an area of anatomy covered by a lesion, the system might call up the area named by the student from the atlas and highlight it for comparison alongside the case image.

5.1.8 Applications of an Atlas-Tutor

A combined Atlas-Tutor offers the possibility of new types of teaching across the range from initial medical education to specialist support. It can take into

account different levels of radiological knowledge through a dynamic model of the student's skills and understanding. Some indications of how such a system might be used are given below.

- a) A general medical student could call on the system for basic instruction in anatomy of the brain, along with a general introduction to MR neuro imaging through presentations of case images illustrating different diseases overlaid with annotated anatomy.
- b) A registrar beginning a specialist training in radiology could be taught a structured approach to reporting, based around a consistent terminology for describing abnormal images. This is similar to the current MR Tutor, but with the atlas providing guidance on the precise meaning of anatomical terms by indicating the relevant structures on atlas plates. It can also support the Tutor in remedial teaching and in managing the selection of cases for training.
- c) A specialist in radiology could be supported in reporting cases through an integrated radiologists' workbench that offers: a library of cases described in a standard terminology; pictorial and symbolic representations of the brain; a unified method of retrieving information on cases, atlas plates and diagnoses; access to online literature; statistical information to support reasoning under uncertainty; and decision aids such as warping models to support the interpretation of images. It enables the specialist to make active connections between situated cases and general medical knowledge in an efficient way.

For all these applications, practical problems have to be overcome. We indicate below some of the issues that need to be addressed in the design of an effective Atlas-Tutor.

5.2 Implementation Issues

Combining a knowledge-based brain atlas with a training system for radiology is a complex task, involving the formalization of radiological knowledge, educational systems design and software engineering. In this section, we discuss some issues relating to the integration of the two systems: a unified framework for knowledge, spatial reasoning techniques and the use of warping models to annotate case images.

5.2.1 A unified framework for knowledge

The MR Tutor and Atlas both contain structured representations of anatomical knowledge. The Atlas team has concentrated on terminology to describe normal anatomy, and the MEDiate team on a language for describing abnormal appearance. These would need to be reconciled, to cover normal and abnormal anatomy and appearance, with each entity given a standard term and placed within a classification system that describes the class hierarchy ('is a'), composition hierarchies ('is part of', 'containment', 'consists of'), spatial relationships ('adjoining', 'overlaps', 'below', 'above and left'), attributes ('shape', 'size' etc.) and links between representational systems (atlas plates, case images, clinical data etc.). The knowledge structure must be anatomically accurate and appropriate for radiological interpretation. It should distinguish between abnormalities that can be seen on an image, and those that are associated with the case (they have been detected on another image or by other methods) but are not visible. It also needs to be computer-accessible, so that it can be referenced by the training system for remedial teaching. The GALEN representation and integration language (GRAIL) [67] might provide a basis for this knowledge representation, though its expressiveness would need to be greatly extended.

5.2.2 Registration of atlas plates on case images

For case-based training, the system needs a method of annotating the case image with the borders and labels of normal and abnormal anatomy. To do this by hand for every image is time-consuming and we are examining methods of automating, or at least assisting, this process.

There are a number of semi-automatic techniques for warping atlas plates into register on case images, and Collins [47] describes an algorithm for fully automated warping and registration. In essence his approach is to provide a set of reference atlas plates in the form of sets of MR slices through a normal brain, each of which has been marked up with the names and borders of its anatomical structures. The algorithm compares, in 3D, volumes taken from equivalent positions within the case images and the reference images. Where the algorithm detects a difference in voxel intensity, indicating a difference in anatomy, it distorts the atlas image into correspondence with the case one. Through a series of iterative steps, first global and then local volumes of the reference image are pulled into register. At present, the Collins algorithm is restricted to normal images and cannot produce accurate registration of peripheral areas such as the gyri and sulci where there is large variation in appearance across cases. The algorithm might be improved by supplying a range of reference images to cover qualitative variations in anatomy.

Although Collins has demonstrated that this algorithm can be applied to a range of normal MR images, extending it to image sequences containing abnormalities is not straightforward. In addition to the general problem of selecting and registering an atlas plate with a case image, an lesion may either invade or displace the surrounding anatomy, producing both changes in intensity and distortion to the anatomical structure. The Collins algorithm might be extended with an additional procedure that distorts the region of the atlas

lattice surrounding the lesion, in a way that depends on both the size of the lesion and whether the lesion is known to invade or compress normal anatomy. This would require intervention by an expert radiologist to assess the type of lesion and in some cases to complete the distortion of the lattice to ensure an accurate registration. Crafting part of the registration by hand would not be as labour-intensive as annotating an entire case. We are not aware of any such implementation, though it would appear to be a tractable problem.

5.2.3 Spatial reasoning

Spatial reasoning can play two different roles in a combined Atlas-Tutor training system: explanation and indexing. If a trainee fails to detect the abnormality or misunderstands its extent (i.e. lacks diagnostic knowledge), or knows the extent of the abnormality but mistakes the anatomy it covers (i.e. lacks anatomical knowledge), a remedial teaching action could show the extent of the abnormality and overlay the relevant atlas plates on the case image. The system could also give written descriptions of the spatial relationships between anatomical structures, or between anatomy and abnormalities. Conversely, the written descriptions could be used as index terms for a query to retrieve abnormalities (e.g. "find all lesions in the left hemisphere within cerebral white but not cortical grey matter").

To provide an explanation or respond to a query, the knowledge base would need spatial reasoning capabilities based on, for example adjacency, consecutiveness, inclusion, and orientation of visible structures [68]. These can be combined into higher-level relations by operators of composition, disjunction and conjunction [69-71].

The time-consuming task of marking up each image with spatial relations could be automated providing that normal and abnormal structures can be outlined and labelled by a warping algorithm. It would need not only to identify

the extent of all diagnostically significant structures, but also to generate symbolic descriptions of their shape and relative positions. A compromise would need be found between coarse and fine detail, and between precision and general applicability.

6. Conclusions

This paper has set out a new approach to computer-based training in radiology that combines a knowledge-based tutor with a medical atlas. There is no guarantee that such a system will be accepted by radiologists as a supplement to human teaching. However, each team has taken a user-centred approach to software design, with each prototype undergoing extensive testing both for usability and educational effectiveness. For the MR Tutor the MEDiate team is carrying out a series of workplace studies to determine current practice in radiology training and to investigate how a computer-based training can fit into the working patterns of trainee radiologists [3]. This will inform the design of future versions.

An Atlas-Tutor offers the promise of combining training and diagnostic support within a single package, providing training on demand through a user-friendly multimedia system, and teaching a systematic approach to radiological reporting, supplemented by extensive atlas-based tools for defining and identifying anatomical structure. This can be done across the spectrum from initial medical education to specialist support.

Designing an Atlas-Tutor training system requires further investigation of knowledge representation, warping models and human computer interaction. It is necessary to design teaching strategies, dynamic student models, and tools for spatial reasoning. Conceptual knowledge of the appearance of abnormal images, in terms of lesion types, shapes, geometric qualifiers, intensities and associated signs needs to be acquired through techniques of knowledge

engineering. Then this knowledge must be supplemented by biomedical knowledge. At present warping models are only suitable for case images of normal brains, but we believe that they could be adapted for images exhibiting abnormalities, as a step towards partially automated interpretation of MR images. These sources of knowledge need to be combined to determine appropriate teaching actions, according to the learner's knowledge and context. This must all be incorporated in an adaptive hypermedia system that can determine the relevant information, computations and teaching actions and then communicate with the user by means of a teaching context.

Acknowledgements

The authors would like to thank C. Barillot, B. du Boulay, G. H. du Boulay, B. Gibaud, N. P. Jeffery, B. A. Teather and D. Teather for their valuable comments and support. This research between France and the UK is funded by the British Council Alliance scheme. The MEDiate Project is supported by a grant from the Cognitive Engineering Initiative of the UK Economic and social Research Council, project number L127251035. The Atlas project was supported by the council of Brittany ("Conseil Régional de Bretagne").

References

- [1] A. Towle, Critical thinking: the future of undergraduate medical education, King's Fund Centre, A Study by the King's Fund Centre in Collaboration with St Bartholomew's Hospital Medical Centre, London, 1991.
- [2] M. Sharples, B. du Boulay, D. Teather, B. A. Teather, N. Jeffery, and G. H. du Boulay, The MR Tutor: computer-based training and professional practice, World Conference on Artificial Intelligence and Education (AI-ED '95), Washington DC, (1995), 429-436.

- [3] M. Sharples, N. Jeffery, D. Teather, B. Teather, and G. du Boulay, A socio-cognitive engineering approach to the development of a knowledge-based training system for neuroradiology, in: B. du Boulay and R. Mizoguchi, eds., *Artificial Intelligence in Education (AI-ED '97)*, IOS Press, Kobe, Japan, (1997), 402–409.
- [4] E. Montabord, B. Gibaud, C. Barillot, S. Garlatti, and I. Kanellos, An hypermedia system to manage anatomical knowledge about brain, presented at *Computer Assisted Radiology*, Berlin, 1993.
- [5] A. Lesgold, H. Robinson, P. Feltovitch, R. Glaser, D. Klopfer, and Y. Wang, Expertise in a complex skill: diagnosing x-ray pictures, in: M. Chi, R. Glaser, and M. J. Farr, eds., *The Nature of Expertise*, Lawrence Erlbaum Associates, Hillsdale, N.J., (1988), 311-342.
- [6] E. Rogers, VIA-RAD: a blackboard-based system for diagnostic radiology, *Artificial Intelligence in Medicine* 7, (1995) 343-360.
- [7] H. P. A. Boshuizen and H. G. Schmidt, On the role of biomedical knowledge in clinical reasoning by experts, intermediates and novices, *Cognitive Science* 16, (1993) 153-184.
- [8] M. Arya, W. Cody, C. Faloutsos, J. Richardson, and A. W. Toga, QBISM: a prototype 3D medical image data base system, *Data Engineering* 16, (1993) 38-42.
- [9] K. H. Höhne, M. Bomans, M. Riemer, R. Schubert, U. Tiede, and W. Lierse, A volume based anatomical atlas, *IEEE Comp. Graphics & Appl.* 12, (1992) 72-78.
- [10] W. L. Nowinski, A. Fang, S. Fang, B. T. Nguyen, R. N. Bryan, and R. Raghavan, Three-dimensional electronic atlas of human cerebral deep structures, presented at *Information processing in Medical Imaging*, Brest, 1995.

- [11] A. Pommert, R. Schubert, M. Riemer, T. Schiemann, U. Tiede, and K. H. Höhne, Symbolic modeling of human anatomy for visualisation and simulation, presented at Visualization in Biomedical Computing, Rochester, Minesota, 1994.
- [12] S. G. Shaw, R. Azevedo, and P. M. Bret, The Effectiveness of hypermedia instructional modules for radiology residents, Canadian Journal of Educational Communication 24, (1995) 245-268.
- [13] R. T. Macura, K. J. Macura, V. E. Toro, E. F. Binet, J. H. Trueblood, and K. Ji, Computerized case-based instructional system for computed tomography and magnetic resonance imaging of brain tumors, Investigative Radiology 29, (1994) 497-506.
- [14] J. F. Brinkley, K. R. Eno, J. W. Sundsten, D. M. Conley, and C. Rosse, A distributed framework for distance learning in anatomy: the Digital Anatomist Interactive Atlas, presented at 19th Symposium on Computer Applications in Medical Care, New Orleans, 1995.
- [15] C. North and F. Korn, Browsing Anatomical Image Databases: a case study of the Visible Human, presented at ACM CHI'96, 1996.
- [16] University of Washington, Radiology Webserver,
<http://www.rad.washington.edu/>.
- [17] S. Sinha, U. Sinha, H. Kangarloo, and K. Huang, A PACS-based interactive teaching module of radiologic sciences, American Journal of Roentgenology 159, (1992) 199-205.
- [18] H. I. Goldberg, S. Fell, H. J. Myers, and R. C. Taylor, A Computer assisted interactive radiology learning program, Investigative Radiology 25, (1990) 947-951.
- [19] Critical Concepts, SymBioSys Clinics,
http://www.laketech.com/SC_MN.HTML.

- [20] Marshall School of Medicine, Interactive Patient Home Page,
<http://medicus.marshall.edu/meducus.htm>.
- [21] M. J. Greenacre, Correspondence Analysis in Practice (Academic Press,
London, 1993).
- [22] G. H. du Boulay, B. A. Teather, D. Teather, M. A. Higgott, and N. P. Jeffery,
Standard terminology for MR image description, in: M. Takahashi, Y.
Korogi, and I. Moseley, eds., XV Symposium Neuroradiologicum,
Kumamoto, (1994), 32-34.
- [23] N. P. Jeffery, Computer Assisted Tutoring in Radiology. PhD Thesis, De
Montfort University, Leicester, 1997.
- [24] J. Talairach and P. Tournoux, Co-Planar Stereotactic Atlas Of The Human
Brain (Georg Thieme Verlag, Stuttgart, 1988).
- [25] G. Szikla, G. Bouvier, and T. Hori, Angiography of the Human Cortex
(Springer Verlag, Berlin, Heidelberg, New York, 1977).
- [26] M. Ono, S. Kubik, and C. D. Abernathey, Atlas of the Cerebral Sulci (Georg
Thieme Verlag, Stuttgart, 1991).
- [27] G. Schaltenbrand and W. Wahren, Atlas for Stereotaxy of the Human Brain
(Thieme, Stuttgart, 1977).
- [28] J. C. Mazziotta, A. W. Toga, A. C. Evans, P. Fox, and J. L. Lancaster, A
probabilistic atlas of the human brain: theory and rationale for its
development, Neuroimage, (1995) 89-101.
- [29] C. Barillot, Basic principles of surface and volume rendering techniques to
display 3D medical data, IEEE Engineering in Medicine Imaging 12, (1993)
111-119.
- [30] H. A. Simon, The New Science of Management Decision (Prentice Hall,
New Jersey, 1977).

- [31] G. Citro, G. Banks, and G. Cooper, INKBLOT: a neurological diagnostic decision support system integrating causal and anatomical knowledge, *Artificial Intelligence in Medicine* 10, (1997) 257-267.
- [32] P. Taylor, J. Fox, and A. Todd-Pokropek, A model for integrating image processing into decision aids for diagnostic radiology, *Artificial Intelligence in Medicine* 9, (1997) 205-225.
- [33] P. Taylor, Computer aids for decision-making in diagnostic radiology — a literature review, *British Journal of Radiology* 68, (1995) 945-957.
- [34] S. Garlatti, I. Kanellos, E. Montabord, B. Gibaud, and C. Barillot, YAKA : un système d'information intensionnel sur le cerveau humain, presented at *Informatique des Organisations et des systèmes d'information et de décision*, Aix en provence, 1994.
- [35] M. Staemmler, E. Claridge, and J. Cornelis, SAMMIE — Software applied to multimodal images and education, presented at *IMAC 93*, Berlin, 1993.
- [36] U. Tiede, T. Schiemann, and K. H. Höhne, Visualizing the Visible Human, *IEEE Computer Graphics and Applications*, (1996) 7-9.
- [37] M. Arya, W. Cody, C. Faloutsos, J. Richardson, and A. W. Toga, A 3D medical image database management system, *Com. Med. Image & Graph.*, (1995) 210-214.
- [38] A. I. Cohn, P. L. Miller, and P. R. Fisher, Knowledge-based radiologic image retrieval using axes of clinical relevance, *Comp. Biomed. Res.* 23, (1990) 199-221.
- [39] R. J. Seitz, C. Bohm, T. Greitz, P. E. Roland, L. Eriksson, G. Blomqvist, G. Rosenqvist, and B. Nordell, Accuracy and precision of the computerized brain atlas programme for localization and quantification in Positron Emission Tomography, *Journal of Cerebral Blood Flow and Metabolism*, (1990) 443-457.

- [40] P. T. Fox, J. S. Perlmutter, and M. E. Raichle, A Stereotactic method of anatomical localization for Positron Emission Tomography, *J. of Computer Assisted Tomography* 9, (1985) 141-153.
- [41] A. C. Evans, T. S. Marrett, J. Torrescorzo, S. Ku, and D. L. Collins, MRI-PET correlative analysis using a volume of interest (VOI) atlas, *J. of Cerebral Blood Flow and Metabolism* 11, (1991) 69-78.
- [42] K. J. Friston, R. E. Passingham, J. G. Nutt, and e. al., Localization in PET images: direct fitting of the intercommissural (AC-PC) line, *J. Cereb. Blood Flow Metab* 9, (1989) 141-153.
- [43] C. Bohm, T. Greitz, R. Seitz, and L. Eriksson, Specification and selection of regions of interest (ROIs) in a computerized atlas, *J. Cereb. Blood Flow Metab.*, (1991) 64-68.
- [44] C. Barillot, D. Lemoine, L. L. Briquer, F. Lachmann, and B. Gibaud, Data fusion in medical imaging: merging multimodal and multipatient image, identification of structures and 3D Display aspects, *European Journal of Radiology* 17, (1993) 22-27.
- [45] C. Barillot, J. C. Gee, L. L. briquer, and G. L. Gouailher, Fusion intra et inter individus en imagerie médicale appliquée à la modélisation anatomique du cerveau humain, *Traitement de Signal* 11, (1994) 513-523.
- [46] K. Natarajan, M. G. Cawley, and J. A. Newell, A knowledge based system paradigm for automatic interpretation of CT scan, *Medical Informatics* 16, (1991) 167-181.
- [47] D. L. Collins, C. J. Holmes, T. M. Peters, and A.C., Automatic 3-D Model-based neuroanatomical segmentation, *Human Brain Mapping* 3, (1995) 190-208.
- [48] J. Niggemann, Analysis and representation of anatomical knowledge, *Applied Art. Intell.*, (1990) 309-336.

- [49] W. L. Nowinski, A. Fang, and B. T. Nguyen, Scaltenbrand-Wharen/Talairach-Tournoux brain atlas registration, presented at SPIE Medical Imaging, San Diego, 1995.
- [50] J. C. Mazziotta, A. W. Toga, A. C. Evans, and P. Fox, A Probabilistic reference system for the human brain, Application to the Human Brain Project: Phase I, June, 1993.
- [51] F. E. Bloom, Databases of brain information, in: A. W. Toga, ed. Three-Dimensional Neuroimaging, Raven Press, New York, (1990).
- [52] H. A. Swett and P. L. Miller, ICON: A Computer-based approach to differential diagnosis in radiology, *Radiology* 163, (1987) 555-558.
- [53] R. Ducournau, Y3, version 3., Sema-Group, rapport technique, mai, 1988.
- [54] C. Barillot, B. Gibaud, E. Montabord, S. Garlatti, N. Gauthier, and I. Kanellos, An information system to manage anatomical knowledge and image data about the brain, presented at Visualisation in Biomedical Computing, Rochester, 1994.
- [55] S. Garlatti, E. Montabord, B. Gibaud, and C. Barillot, A methodological approach for object knowledge bases, presented at Expert Systems 95, Cambridge, 1995.
- [56] D. L. Collins, G. L. Goualher, A. C. Evans, and C. Barillot, Cortical constraints for non-linear cortical registration, in: K. H. Höne, ed. Lecture notes in Computer Sciences: Visualization in Biomedical Computing, Springer Verlag, Berlin-Heidelberg, (1994).
- [57] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans, Automatic 3D inter-subject registration of MR volumetric data in Standardized Talairach Space, *J. Computer assist. Tomog.* 18, (1994) 192-205.

- [58] J. C. Gee, C. Barillot, L. L. Briquer, D. R. Haynor, and R. Bacjy, Matching structural images of the human brain using statistical and geometrical image features, presented at SPIE, Visualization in Biomedical Computing, 1994.
- [59] J. C. Gee, L. L. Briquer, C. Barillot, D. R. Haynor, and R. Bacjy, Bayesian approach to the brain image matching problem, SPIE Medical Imaging: Image Processing, vol. 2432, SPIE, (1995), 154-156.
- [60] J. C. Gee, L. L. Briquer, C. Barillot, and D. R. Haynor, Probabilistic matching of brain images, in: Y. Bizais, C. Barillot, and R. Dipaola, eds., Computational Imaging and Vision: Information Processing in Medical Imaging, vol. 2432, Kluwer Academic Publishers, (1995), 113-126.
- [61] M. Bieber, Providing information systems with full hypermedia functionality, presented at IEEE, 1993.
- [62] M. Bieber, On integrating hypermedia into decision support and other information systems, Decision Support Systems 14, (1995) 251-267.
- [63] E. Montabord, S. Garlatti, B. Gibaud, and C. Barillot, Interaction between a hypermedia and a knowledge base management system about the brain, in: E. C. Ifeachor and K. G. Rosen, eds., Neural Networks and Expert Systems in Medicine and Healthcare, University of Plymouth, Plymouth, (1994), 136-145.
- [64] E. Montabord, B. Gibaud, and C. Barillot, HYPER-YAKA : base de connaissances et hypermédia sensible au contexte, Langage et Modèles à Objets, , Nancy, (1995).
- [65] R. W. Howard, Concepts and Schemata (Cassell, London, 1987).
- [66] J. S. Brown, R. R. Burton, and J. d. Kleer, Pedagogical, natural language, and knowledge engineering techniques in SOPHIE I, II and III, in: D. H.

- Sleeman and J. S. Brown, eds., *Intelligent Tutoring Systems*, Academic Press, London, (1982), 227-282.
- [67] A. L. Rector, S. Bechhofer, C. A. Goble, I. Horrocks, W. A. Nowlan, and W. D. Solomon, The GRAIL concept modelling language for medical terminology, *Artificial Intelligence in Medicine* 9, (1997) 139-171.
- [68] O. Salvetti, G. Braccini, R. Evangelista, and M. Freschi, An intelligent system for the diagnosis of complex images, *Artificial Intelligence in Medicine* 8, (1996) 167-185.
- [69] A. G. Cohn, D. A. Randell, Z. Cui, and B. Bennett, Qualitative spatial reasoning and representation, presented at QUARDET, Barcelona, 1993.
- [70] D. Hernandez, *Qualitative Representation of Spatial Knowledge* (Springer Verlag, Berlin, 1994).
- [71] T. P. McNamara, Mental representation of spatial relations, *Cognitive Psychology*, (1986) 87-121.

Figure 1. Screen display of the MR Tutor, showing an overview plot for Glioma (dark points) and Infarct (light points) at the lower right of the screen.

Figure 2. Architecture of the Atlas Decision Support System

Figure 3 Screen display of the Atlas system

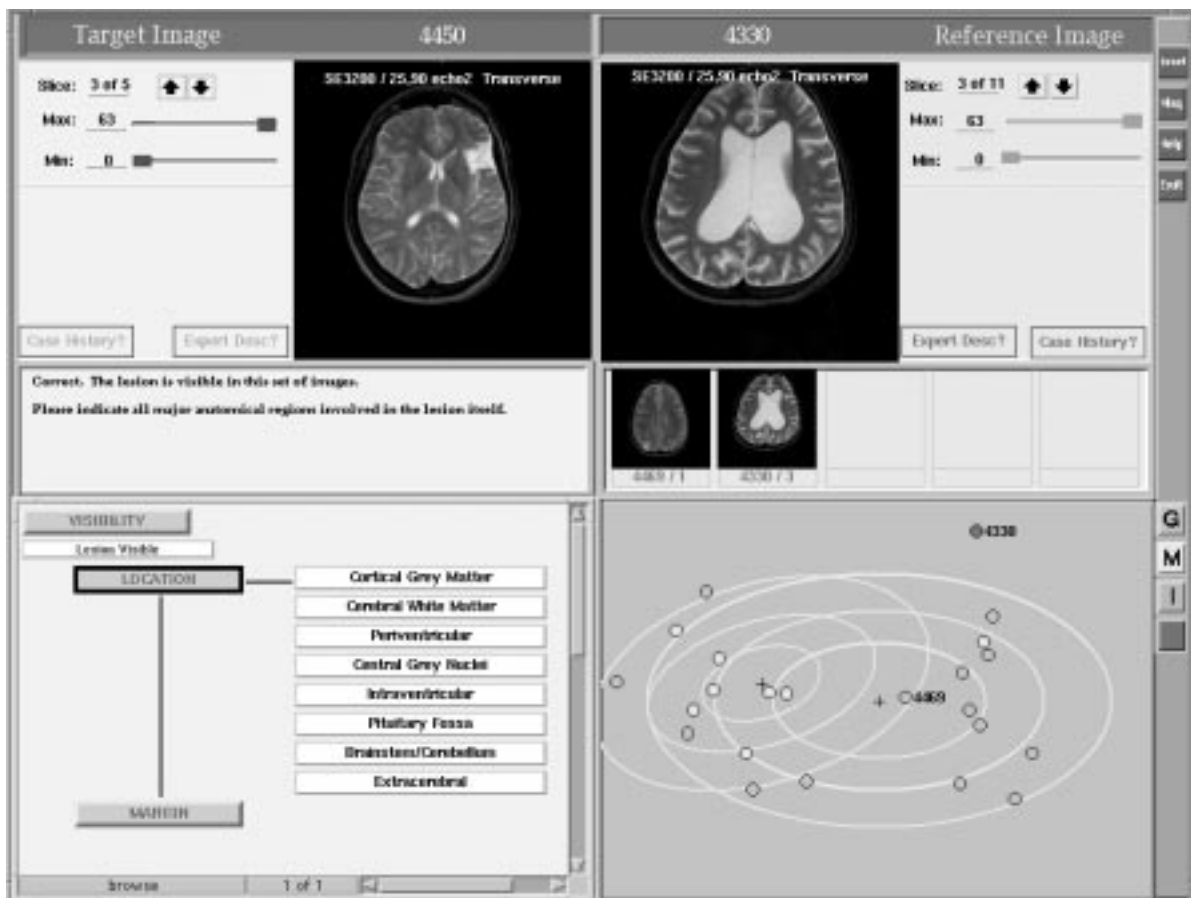


FIGURE 1

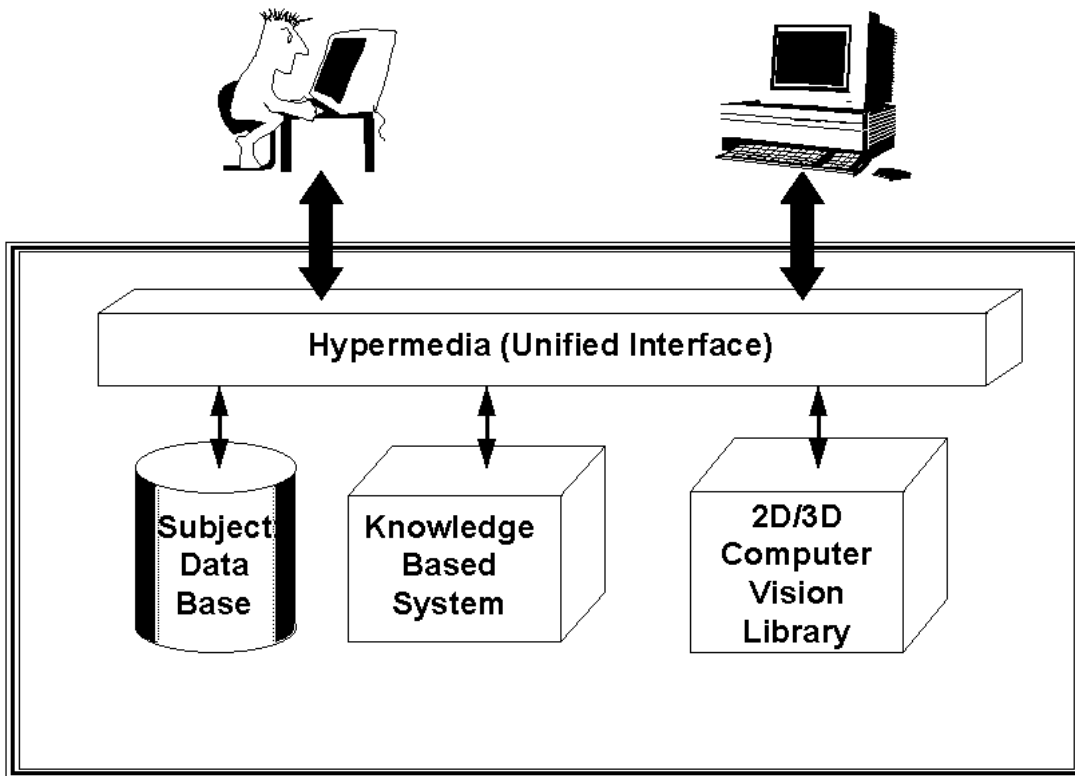


FIGURE 2

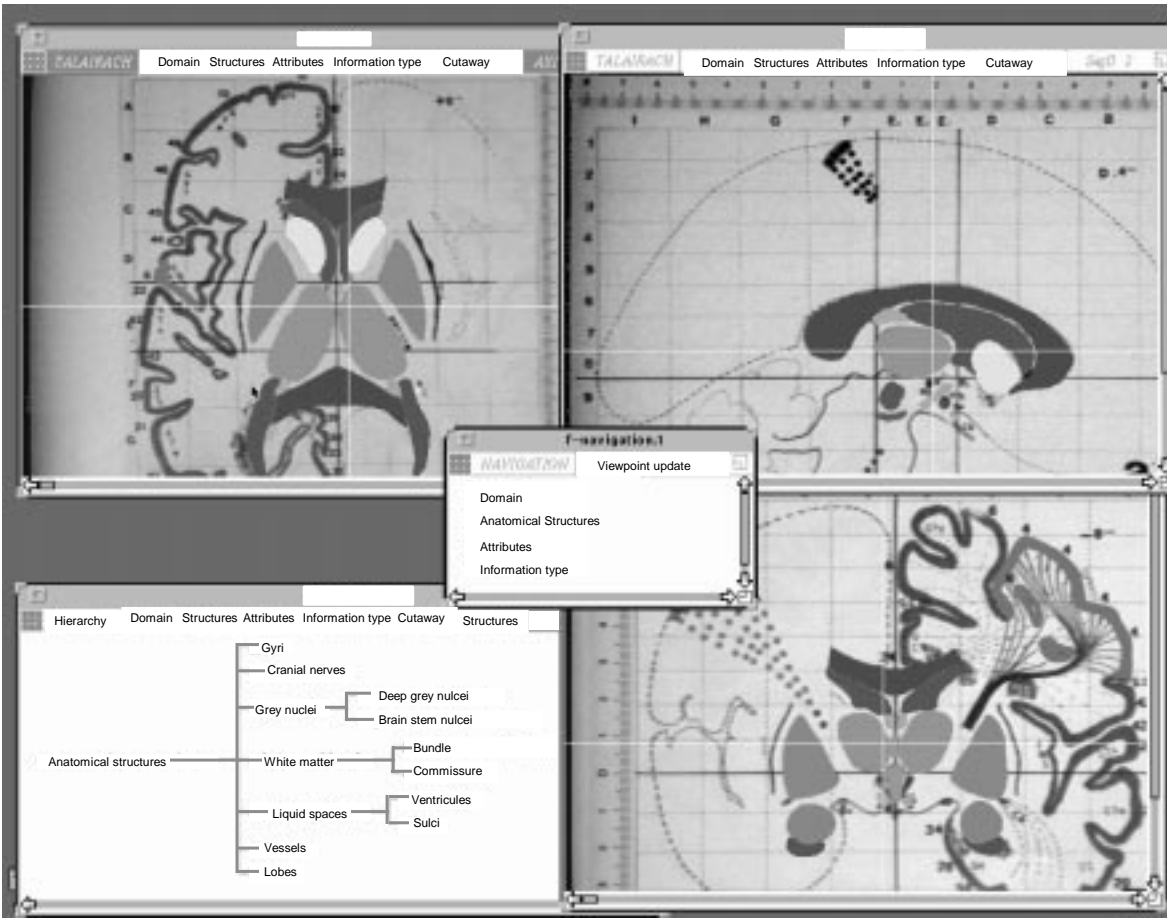


FIGURE 3

