A Socio-Cognitive Engineering Approach to the Development of a Knowledge-based Training System for Neuroradiology

Mike Sharples, Nathan Jeffery,
School of Cognitive and Computing Sciences,
University of Sussex, Falmer, Brighton BN1 9QH, UK.

Derek Teather, Briony Teather,
Medical Systems Research Group, School of Computing Sciences
De Montfort University, Leicester LE1 3BG, UK.

George du Boulay
(Visiting Professor at the Medical Systems Research Group, De Montfort University)
Institute of Neurology, Queen Square, London, WC1N 3BG, UK.

Email: rad@cogs.susx.ac.uk

Tutoring systems could satisfy a demand of many professions for structured case-based training, but to be accepted they need to be robust, authoritative and matched to the needs of trainees in the workplace. This paper outlines a methodology for the development of knowledge-based training that integrates software, task, knowledge and organizational engineering. It consists of a set of "building blocks" that specify the type of activities needed to develop a complete knowledge-based training system, while allowing flexibility in the choice and ordering of specific design techniques. The approach is illustrated by a project to develop the MR Tutor, a knowledge-based training system for neuroradiology. The building blocks for this project have included an analysis of published studies of cognitive processes in medical image interpretation, elicitation and refinement of knowledge from an expert neuroradiologist, workplace studies of radiology training and experiments with new techniques for data visualization. The MR Tutor gives trainee radiologists the experience of observing and analysing a large archive of cases and practice in comparing their interpretations with those of experts. It is based on a structured language for describing abnormal appearance in Magnetic Resonance images of the head, and it uses a novel "overview plot" to visualize and interact with the image archive. The development methodology has been followed to the stage of implementing a robust integrated system.

1. Introduction

Despite many innovations in artificial intelligence and medical education in recent years, including tutors derived from expert systems, case based tutors and training via the Internet, computers are still not widely used for medical training. Our study of UK institutions for radiology training found that none of them offer computer-based training as an integral part of the syllabus. Why has the medical profession been so reluctant to use computers to supplement conventional medical training? Is there an approach to software design that might offer more chance of adoption by medical trainers? This paper addresses these questions, in the context of the MEDIATE project to develop a knowledge-based training and diagnostic support system for neuroradiology (the MR Tutor).

This research is funded by grant L127251035 from the UK Economic and Social Research Council.

2. Development methodology

Very few projects in AI and education have put equal emphasis on software, task, knowledge and organizational engineering. Some systems have demonstrated powerful techniques in knowledge representation, but are virtually unusable. Others have attractive interfaces, but do not address a clear educational need. Yet others meet a well-identified need, but do not fit easily into the classroom or workplace.

Instead of relying on intuition or received wisdom about training needs and solutions we have developed a "building block" methodology for the design of knowledge-based training and cognitive support systems (see Table 1) that integrates contributions from software engineering, task engineering, knowledge engineering and organizational engineering. To construct a successful integrated system requires all the interlocking building blocks, but it is not necessary to build systematically from the bottom up. The development team may work on one "pillar", such as knowledge engineering, up to the stage of system requirements, or they may develop an early prototype based on a task analysis but without a systematic approach to software engineering. Each block specifies one type of design activity, but how that activity is carried out depends on the particular subject domain and training situation.

Maintain	Installed system	New task structure	Augmented knowledge	New organizational structure
Evaluate	Debugging	Usability studies	Conceptual change, Skill development	Organizational change
Integrate	Integrated System			
Implement	Prototypes, Documentation	Interfaces, Cognitive tools	Knowledge representations	Communications, Network resources
Design	Specifications, Algorithms and heuristics	Interface designs, Interaction designs	Domain representations, User models, Support strategy	Interaction design, Communication Support
Interpret	System Requirements			
Analyse	Requirements analysis	Task analysis: goals, objects, methods	Knowledge elicitation: concepts, skills	Workplace studies: practices, interactions
Survey	Existing systems	Conventional task structure	Domain knowledge	Organizational structure
	Software Engineering	*Task Engineering*	*Knowledge Engineering*	*Organizational Engineering*

Table 1. "Building block" methodology for socio-cognitive system design

3. Overview of the MR Tutor

The remainder of the paper describes the MR Tutor and outlines how the MEDIATE project has followed the "building block" method in developing the system. Each section of the paper covers one pillar of the method, describing the methods and outcomes. Companion papers [1][2] cover aspects of the system design in more detail.

Figure 1 shows the interface to the current MR Tutor. The purpose of the system is to supplement professional training in neuroradiology by offering computer-based tutoring based on a structured case archive. The main features of the system are described here in brief, to provide an introduction to its design and function.

The screen is divided into two halves. The right side of the screen allows the trainee to interact with the image archive. The Tutor provides an visual overview of a pathology, as an "overview plot" shown on the bottom right of the display. The trainee can select a case to examine by clicking on a point which is then lit up on the overview plot. The associated stack of images then appears in the pane at the top right of the screen. A control panel enables the trainee to move up and down the slices and to "window" an image by adjusting its grey levels. The trainee can then call up similar cases, and more or less typical ones, by selecting other appropriate points on the plot. Slices can be dragged to the "gallery" below the main image pane. The tutor also provides a magnifier, as a separate window showing an expanded section of an image.

The left half of the screen is the interface to the tutoring component. Either the trainee or the system can select a case for tutoring (a "target case") and this is indicated on the overview. The target images can be viewed and windowed as for the reference images, and the trainee can make visual comparisons between the target case and any reference cases. The trainee then describes the target case by selecting terms from pop-up menus of image features and their descriptors. The system gives a short confirming or remedial message each time the trainee completes a section of the structured description. A typical response is shown in Figure 1, in the pane above the menus.

When the description is complete, the Tutor shows the expert's description of the case alongside the trainee's in the menu window. It also generates a point on the overview plot corresponding to the trainee's description, so that the trainee can make comparisons with the target point indicating the expert's description of the case. It shows, for example, whether the trainee's description is more or less typical than the expert description, or if it is close to other cases.

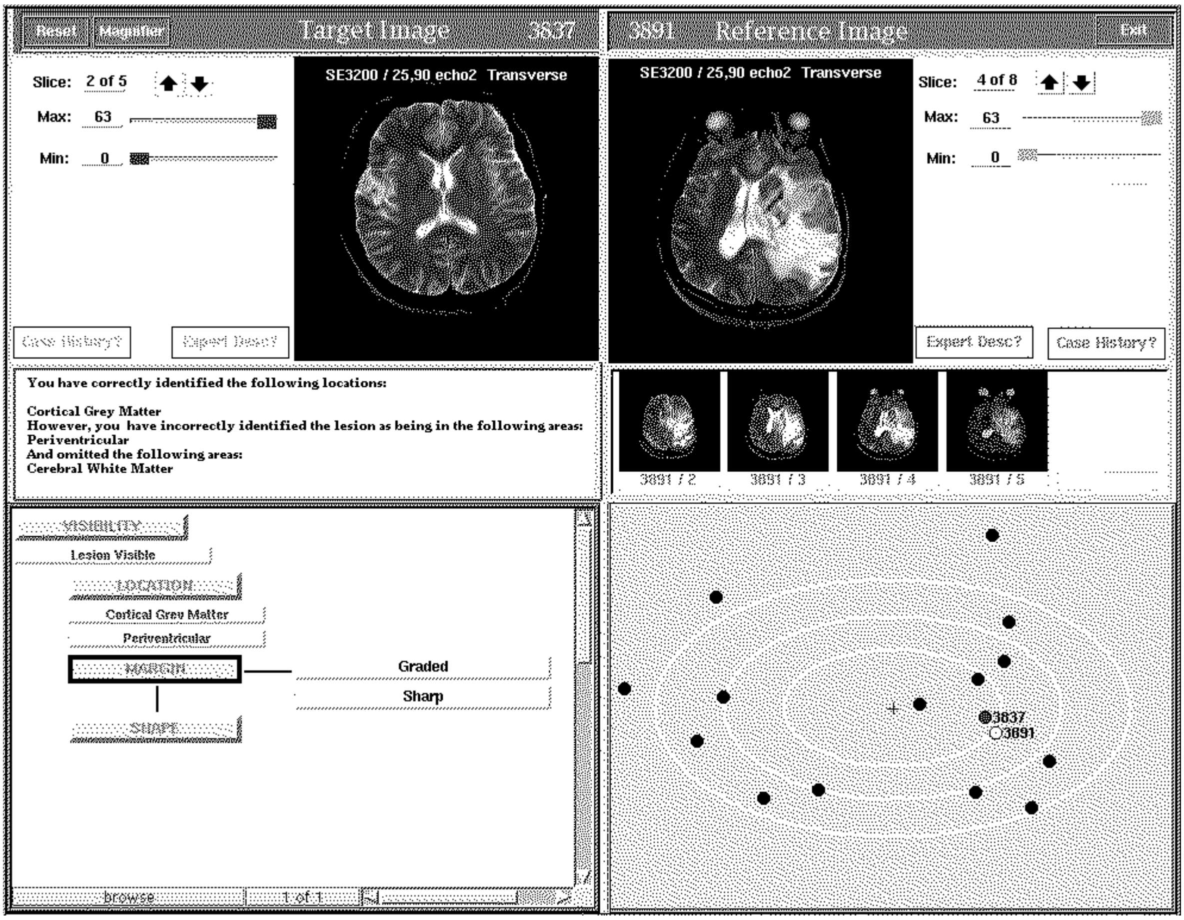

Figure 1. The interface to the current version of the MR Tutor

4. Software Engineering

4.1 Software Survey and Analysis

We can describe existing computer-based training systems for radiology in terms of a spectrum from teacher-led to learner-directed, with the main types being: computer-assisted instruction, simulations, case-based tutoring and diagnostic support, and reference aids.

An example of computer-assisted instruction is MITS, the Medical Image Teaching System [3]. A typical MITS lesson presents one or more images, a question asking the trainee to identify abnormalities, and a list of possible answers. A computer-assisted instruction system is limited by its inability to know what it is teaching. It cannot engage the student in a discussion of the images beyond what has been anticipated by the programmer, nor can it track a student's developing expertise, to respond in a style and detail that matches the learner's current understanding.

Reference aids overcome some of the limitations of computer-assisted instruction by providing passive resources rather than active teaching. Typically they allow the trainee to browse through a library of images indexed by pathology or by a free text description. They offer a useful adjunct to human or textbook teaching, but the annotations and description associated with the images do not form a systematic representation of knowledge, that could be used to answer complex student queries or provide a knowledge base for computer-assisted training.

Simulations address other pedagogic limitations of computer-assisted instruction, 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 MRI scanner, to a virtual patient on which the student can perform tests or surgery. Simulations can be powerful aids to learning, but a simulation alone does not provide teaching or direct cognitive support.

Few active knowledge-based tutors have been developed. The CT Brain Tutor developed at the Medical College of Georgia [4] trains radiology residents to diagnose brain tumours from CT and MRI scans. It is composed of three modules: Case Retrieval, Teach Me, and Atlas. The Case Retrieval module presents images from a library of 120 cases indexed by case history and radiological features. The indicative features are pathology based ('tumour', 'oedema', 'calcification') and do not form a comprehensive representation of knowledge about images. Teach Me is an educational program to assist a trainee in memorising the presented tumour pattern. It has a series of windows displaying a section of normal brain, brain with tumour, coloured features of tumour superimposed on the pathological images, a textual description of the features and a colour key. The Atlas module displays nine sections of normal brain, with detailed explanations of visible anatomical structures. The CT Brain Tutor shows the potential of active knowledge-based tutoring, but it has a limited knowledge representation and range of tutoring strategies.

The Radiology Tutor [5] generated Socratic tutorials about the appearance of chest X-rays. It was able to conduct a sophisticated dialogue with a trainee in natural language but was constrained by its limited natural language understanding, being only able to parse a limited subset of English.

RUI [6] is a generalisation of the Radiology Tutor into a general-purpose authoring shell. Its Specification Tool enables a domain expert along with a programmer to enter knowledge of an imaging modality. The Image Description Tool can be used by the radiologist alone to annotate and describe a database of images. The Learning Tool calls on the knowledge-base and annotated images to automatically generate tutorial dialogues. Demonstration tutors have been implemented in RUI for MR, CT, Ultrasound and X-ray.

RUI and the Radiology Tutor take a conventional Intelligent Tutoring Systems approach to training, with the system acting as a simulated human teacher. An initial analysis of requirements for the MR Tutor has shown that the prime need is not for a tutor with deep knowledge but limited scope for learning. What trainees most require is an extensive archive of well-indexed cases, with powerful tools for browsing and image visualisation, accompanied by an unobtrusive tutor.

4.2 Software Design and Implementation

Thus, we have concentrated on providing tools for a trainee to visualise the spread of cases within a pathology, to compare pathologies, and to access and compare cases by their similarity and typicality of appearance. The current MR Tutor offers a 'lightweight' tutor. The trainee can choose any case for tutoring, can suspend the teaching at any time to move to another case, and can interleave teaching with browsing through the archive.

Our general approach to software engineering has been to construct a series of demonstrator systems, including a mock-up in SuperCard, a fully functioning system in HipWorks (a multimedia development environment under UNIX) and restricted versions to test innovations, such as a rule-based tutor, and demonstrations of interfaces for displaying overviews of the cases. This cycle of iterative 'off-line system development' ensures that a number of new designs can be thoroughly tested in parallel before they are incorporated into the main system.

We have adopted, from the Writer's Assistant project, the approach of 'provable need'. Any new addition to the system interface or functionality must be shown to address some explicit need of radiology trainers or trainees. Thus, an image 'magnifying lens' was only added when it became clear from trials that this was demanded by the radiologists. The aim to produce a system with an uncluttered interface and functions that meet the needs of training rather than just demonstrate programming skills.

5. Task Engineering

5.1 Task Survey and Analysis

We have not yet carried out a thorough task analysis of radiological interpretation, but we have drawn on studies conducted by Lesgold and colleagues [7]. They observed novice and expert radiologists as they examined and marked up X-rays while thinking out loud. Lesgold et al. propose that radiologists carry out a multi-stage process of interpretation. On first seeing a film they automatically invoke a mental schema that covers the salient features, resulting in one or more tentative diagnoses. This triggers a cognitive process of active search for other cues in the image along with case data and biomedical knowledge that might constrain the interpretation. Lastly, they articulate their findings as a verbal report.

Azevedo and Lajoie [8] have studied the teaching of radiology in Canada by observing and videotaping a teaching round involving a staff member and residents at Montreal General Hospital. They found that the radiologist typically "walked" a junior resident through the interpretation process. For more competent residents, the radiologist "scaffolded" their interpretation process by providing hints and directions as to where to direct attention. On occasion the radiologist articulated her reasoning process beginning with assignment of probability to pathological features, followed by the systematic elimination of differential diagnoses, until she reached a definitive diagnosis. This account of rapid perceptual analysis followed by hypothesis-driven reasoning to reject and confirm competing diagnoses fits with Lesgold's account of the reasoning processes of an expert radiologist.

These studies have some limitations as a foundation for the design of the MR Tutor. They were conducted with chest X-rays and the findings may not transfer over to the more complex area of MR imaging of the head. They do, however, match the descriptions of the process of interpretation and training provided by our own informal discussions with consultant radiologists.

5.2 Task Design and Implementation

The implication of the task studies for system design is that a computer-based tutor for radiology should support the multi-stage process of interpretation, by allowing a trainee to make a rapid initial interpretation and then assisting in the process of searching for other salient cues in the image. The system should also scaffold the process of interpretation by offering hints and directions, and should be able to articulate the stages of diagnosis.

The current system supports the multi-stage process of interpretation by enabling the trainee to scan rapidly through the set of images associated with the target case, and to make rapid visual comparisons between target images and ones from the archive. Then, the system guides the trainee through a structured description of the case, with the system recording each student action and giving tutorial responses. It also provides extensive scaffolding. A context-dependent help system offers a definition or illustration of terms used by the system, to help the trainee to enter a full and accurate description. As the trainee enters a description of a target case, by moving the mouse over a menu of descriptors (e.g. the terms for "major location" shown in Figure 1), the system "lights up" the points on the overview plot corresponding to cases in which the feature is present. This shows at a glance the position of the case being tutored relative to those with a particular feature. An aim for the future is to gradually remove the scaffolding as the system detects that the trainee becomes more competent, but this aspect of the system has not yet been implemented.

The main innovation in interface design is the "overview plot". The statistical technique of multiple correspondence analysis is applied to a structured description of each case to produce a 2-

dimensional plot showing the distribution of cases for a pathology, such that typicality and similarity can be read directly from the display. The nearer a case is to the centre of the display, the more it is typical of the pathology (a function of the number of features it shares with other cases). The similarity of any two cases is shown by the distance between their two points. The trainee can call up an individual case and view its stack of images by clicking on its marker point. Two or more pathologies can be overlaid on the same plot, giving a powerful way of indicating "difficult to diagnose" cases on the border between pathologies.

The methods for deriving, displaying and interacting with the overview plot are quite general and can be applied to any binary data forming a multidimensional space. For example, it could be used as a general method for visualising and interacting with a picture library, a database of topics for self-study, or a concept map, providing that each picture, topic or concept has been marked up in advance with a structured description of its content.

6. Knowledge Engineering

6.1. Knowledge Survey and Analysis

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. Although standard terminologies are widely used for disease categories and anatomical structures, 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 of describing abnormalities, there is widespread agreement on the need to develop a more structured approach to reporting, so that radiologists can exchange findings in an agreed language using terms that have been precisely defined. The value of a standardised description language for images has been widely recognised in the literature of clinical radiology.

6.2 Knowledge Design and Implementation

A collaborative research project between the Medical Systems Group, De Montfort University, and the Institute of Neurology, London, has developed and validated an image description language (IDL) for MR images of the head for a wide range of image sequences. The language describes the appearance of the images rather than the underlying pathology (although the appearance is influenced by the pathology). The terms are understandable to neuroradiologists (for example, "overall appearance of lesion, homogeneous", "lesion margin, graded", "location, cerebral white matter") and are supported by precise definitions or, in the case of subjective terms such as shape, by indicative examples.

The IDL offers a sound basis for indexing the cases, and the MEDIATE project now has an archive of some 1200 cases fully described using the terms of the language. The array of feature descriptors for each case provides a knowledge base both for the tutorial responses and for automatic construction of the overview plot. The IDL defines a multi-dimensional feature space, where each dimension represents one feature value (e.g. "lesion size large") with two states, present or absent. The description for each case represents a single point in the feature space. To compute the overview plot, the multiple correspondence analysis algorithm projects this multi-dimensional feature space down onto a two-dimensional surface, such that the data points indicating the cases are maximally spread out over the surface and that the plot gives a direct indication of the typicality and similarity of cases.

The domain knowledge for the IDL has been elicited from a single expert, who has also described each case in the archive. Although eliciting knowledge from a single expert is not recommended as a basis for developing expert systems, we believe this is an appropriate method for radiology training where there is no standard approach to reporting and where each radiologist has a different terminology. The most important requirement is to give an operational definition of each term and to validate the IDL, so that the trainee is acquiring a precise and comprehensive terminology.

7. Organizational Engineering

7.1. Organizational Survey and Analysis

We have taken a four-level approach to conducting workplace studies, with each level illuminating the practice of radiology training, and also giving a framework and issues to be addressed for the next level. Level one, which is now complete, involved collecting course material from institutions such as the Royal College of Radiologists and carrying out telephone interviews with radiology lecturers, to gain an overview of radiology training throughout the UK. Level two is a study of teaching practice through observation of teaching sessions. Level three will consist of interviews with radiology trainers and trainees, to uncover their conceptions of the work and to discover issues, such as teaching problems and current use of technology, that can indicate how a computer-based training system could be situated within current training practice. For the fourth level we shall carry out heuristic evaluations of versions of the MR Tutor with radiology staff and trainees as it is revised for use in everyday training. The findings reported in this section are from the level one study.

New technology forms an integral part of modern medicine. In radiology, the new modalities of Computerised Tomography (CT) and Magnetic Resonance Imaging (MRI) capture images digitally, under computer control. The institutions we surveyed all provide personal computers for the use of registrars, and all have purchased CDROMS or laser disk reference systems with databases of medical images. We found no general hostility towards computer technology. There are, however, specific reasons why the radiology profession is reluctant to embrace computer-based training.

First, most consultant radiologists within non-digital departments prefer to interpret from film (rather than film scanned into digitised images) as it provides a high level of detail and allows multiple images to be viewed alongside each other. However, some consultants working in departments with direct digital capture suggest that digital images are better for interpretation as they offer flexibility in adjusting display parameters such as contrast and brightness. Second, radiologists are suspicious of the quality of information provided by computer-based systems. For example, none of the institutions we surveyed use multimedia atlases, in part because these do not yet have the authority of text-based atlases by renowned authors. Third, the style and format of training varies considerably depending on the trainee's level of ability. Lectures and tutorials form a small part of the training of radiologists. Most learning comes through clinical experience and apprenticeship in the reporting of cases. Self study also forms a major part of training on some courses, with one day a week being devoted to self study and a further 2-3 hours a night expected.

7.2 Organizational Design and Implementation

It is clear that an important factor influencing the acceptance of the system is its perceived authority. This governs the entire design. The interface should give the general impression of a robust medical system and it should employ interface objects, such as controls for windowing the images, consistent with those on standard radiology equipment such as MRI consoles. The terminology throughout should be appropriate to the profession of radiology, and also signify the knowledge of a respected authority in the field. The teaching strategy should mesh with the case sessions that form the major part of a radiologist's training.

The most appropriate setting for computer-based training in radiology is as part of self-study. The trainee is working alone without the distraction of case meetings or ward rounds, and has access to supplementary material. This means, however, that the system must be self-contained and available on a personal computer. Any queries from the trainee must either be answered by the system, or by remote access to a human tutor, or the context of the query must be saved so that the trainee can ask a tutor at a later time. We are investigating all three methods of giving supplementary advice.

The great advantage of a personal system is that it can be integrated with other learning and diagnostic tools, to provide continuing support for diagnosis, research and professional development. A consultant radiologist, for example, might use the system to create a personal archive by reporting new cases in the structured language which could then be indexed via the overview plot and compared with the existing cases.

8. System Integration and Evaluation

The individual pillars of software engineering, task engineering, knowledge engineering and organizational engineering underpin the development of a knowledge-based multimedia training system, but they need to be linked together into, first, a set of system requirements and, second, an integrated system. In developing the initial system we have been guided mainly by the need to provide an authoritative representation of knowledge and an elegant and intuitive interface.

The current prototype runs on personal computers under Linux. On a 133MHz Toshiba Tecra laptop the response is almost instantaneous, there is no appreciable delay in retrieving images from disk or in receiving a tutorial response from the system. An initial formative evaluation of the system shows that the main functions can be learned in less than 10 minutes, that radiologists engage with the case tutorials, and that they find the overview plot a powerful aid to learning.

9. Conclusions

There would appear to be no theoretical or practical reasons why computer-based training should not be widely adopted in medicine, but we argue that designers of such systems have consistently underestimated the difficulties. The development of successful knowledge-based multimedia systems for training in the professions does not just involve imaginative programming and extensive knowledge elicitation. The system must be designed to fit into the workplace and the schedule of a busy trainee. It should have the authority of a good textbook, but should supplement rather than replace book learning. It should be based on sound cognitive and social theories of tutoring, skill acquisition and professional development. Most of all, to stand a chance of being accepted as an integral part of professional training it needs to be developed through close collaboration between domain experts, trainers, system designers and instructional designers. In this paper we have set out a general methodology for developing knowledge-based training systems and shown how it has guided the design of a multimedia system for training in neuroradiology.

10. References

[1] M. Sharples, B. du Boulay, D. Teather, B. Teather, N. Jeffery, and G.H. du Boulay, The Cognitive Basis for an MR Image Tutor. *Proceedings of East-West Conference on Computer Technologies in Education, Part 1*, The Crimea, 1994, pp. 214-219.

[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. *Proceedings of World Conference on Artificial Intelligence and Education (AI-ED 95),* Washington DC, 1995 pp. 429-436.

[3] 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.

[4] 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, 4 (1994) 497-506.

[5] M. Sharples, The Radiology Tutor: Computer-based Teaching of Visual Categorisation. In D. Bierman, J. Brueker & J. Sandberg (eds.) *Proceedings of 4th International Conference on AI and Education,* IOS Press, Amsterdam, 1989, pp. 252-259.

[6] A.I. Direne, Methodology and Tools for Designing Concept Tutoring Systems. *Proceedings of World Conference on Artificial Intelligence and Education (AI-ED 93),* Edinburgh, 1993, pp.55-58.

[7] A. Lesgold, H. Rubinson, P. Feltovitch, R. Glaser, D. Klopfer, and Y. Wang, Expertise in a Complex Skill: Diagnosing X-ray Pictures. In: M. Chi, R. Glaser, & M.J. Farr (eds.), *The Nature of Expertise,* Lawrence Erlbaum Associates, Hillsdale, N.J., 1988, pp. 453-494.

[8] R. Azevedo and S.P. Lajoie, Learning Styles Underlying Radiological Expertise. *Project Safari Report 005*. McGill University, 1995.