

Training in Image Description and Diagnosis for MR Radiology of the Brain

Key intellectual contribution

The key intellectual contribution is the design, implementation and testing of a computer-based training and decision support system for neuroradiology that is based on a well-founded Image Description Language and novel method of visualising the distribution of cases within and across diseases.

Previous publications

The work has been extensively published in journals and proceedings of neuroradiology, medical informatics, and artificial intelligence and education. This paper is an updated and abridged version of a paper that was presented at the Joint European conference on Artificial Intelligence in Medicine and Medical Decision Making. It provides a concise summary of work to date on the project.

Related work

The most closely related work is that of Azevedo and colleagues on the design of a computer-based training system for mammography, informed by studies of the cognition and practice of radiology. Some of the techniques employed in our work, such as the use of Multiple Correspondence Analysis to produce visualisations of high-dimensional data, have been employed elsewhere, but their application to computer-based training in radiology is novel, as the development and use in computer-based training of a structured Image Description Language for MR radiology of the brain.

Association with Clinical Practice

The Image Description Language has been developed through a close collaboration over many years with an eminent neuroradiologist. It has evolved through a process of critical appraisal with other senior neuroradiologists. The requirements for the training system are based on a two-year investigation of the practice of neuroradiology reporting and training in two teaching institutions. The overview plot has been tested through a formative evaluation involving seventeen subjects, comprising four novices (with no knowledge of radiology), nine intermediates (4th year medics and radiographers with some knowledge of anatomy and imaging) and four expert neuroradiologists.

Training in Image Description and Diagnosis for MR Radiology of the Brain*

Benedict du Boulay^{a†}, Lisa Cuthbert^e,

Fernando De Andres Garcia^b, George du Boulay^c, Nathan Jeffrey^b, Mike Sharples^d,

Briony Teather^b, Derek Teather^b and Emma Whitcombe^a

^aUniversity of Sussex, ^bDe Montfort University, ^cInstitute of Neurology,

^dBirmingham University, ^eOxford Brookes University

Abstract. We have developed a system that aims to help trainees learn a systematic method of describing MR brain images by means of a structured image description language (IDL). The training system makes use of an archive of cases previously described by an expert neuroradiologist. The system utilises a visualisation method – an Overview Plot – which allows the trainee to access individual cases in the database as well as view the overall distribution of cases within a disease and the relative distribution of different diseases.

1 Introduction

We have developed a training system [10–13] that aims to help trainees to learn how to describe MR brain images in a systematic way by means of a structured image description language (IDL). This language allows clinically meaningful features of MR brain images to be recorded, such as the location, shape, margin and interior structure of lesions. The training system makes use of images from an archive of about 1200 cases, previously described in detail by an expert neuroradiologist.

The image description training system employs a visualisation method – an Overview Plot – which allows the trainee to view and access (i) the images themselves, (ii) the written descriptions of the individual lesions in the image, and (iii) two dimensional representations of the multi-dimensional distribution of all cases of a disease chosen from the archive. The two dimensional representations relate to, and are calculated from, the descriptions of the lesions as seen in each image sequence of a case. Thus one can view the overall distribution of appearance of cases within a disease and the relative distribution of different diseases, one against another. To this extent it is a kind of case-based training system that provides a visual indexing mechanism to cases similar to the case in hand.

Two principles have guided the development of the system so far. First, the system is deliberately aimed to support and train the radiologist's inferences from what can be observed *in the images*. In particular, the two dimensional representations are currently based on lesion appearance and confirmed diagnosis, but not on clinical signs and symptoms. The reconciliation of those inferences with other sources of data, such as the clinical history, is a matter for the user. Second, the design exploits as far as possible radiologists' visual-spatial reasoning rather than simply offering numerical information about diagnostic probabilities.

2 Visual Decision Support

2.1 Image Description Language

The basic domain representation underpinning the system is an archive of cases with confirmed diagnoses, all described by the same expert (G. du Boulay) using the IDL. These include separate descriptions for each image sequence/echo including detailed descriptions (e.g. the *region*, *major position*, *exact location*, *margin*, *structure*, *shape*, *area*, *conformity to anatomical feature*, *interior pattern* (if any), and its *intensity*) of the lesion (or the largest of each type of lesion visible, where there are multiple lesions), as well as *correspondence* between described parts of lesions seen under different sequences and descriptions of *atrophy*, *other signs* and other *abnormal signals* for the case as a whole [3].

For the purposes of the prototype description training system a simplified version of the description language has been used. It provides an initial set of terms to support discussion and sharing of knowledge amongst trainee

* This paper is an updated and abridged version of [2]

† Email: bend@cogs.susx.ac.uk

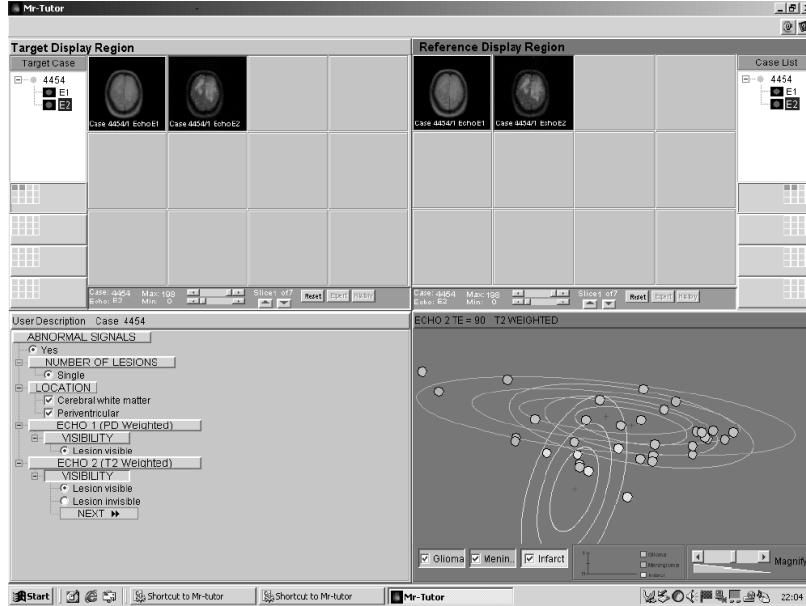


Figure 1. The small world of Glioma, Meningioma and Infarct.

neuroradiologists and their supervisors. It also serves as a structured representation of knowledge for the MR Tutor, enabling it to generate remedial responses to student errors.

2.2 Display of Small Worlds

We can consider each image sequence for a case as occupying a point in a many-dimensioned space of description features. For the current version of the simplified language this space has some 50 dimensions, where each point is a vector of binary values, each representing the presence or absence of a particular feature. Multiple Correspondence Analysis (MCA) is a statistical technique for data reduction and visualisation [6]. It is used here to reduce the dimensionality down to two so as to provide a ready means of overviewing the data. It does this for each image sequence by finding that plane which best spreads out the subset of cases under consideration. MCA is similar to principal components analysis (PCA) but is applied to categorical/binary data as opposed to scalar data and allows for all possible pairwise associations in the data. PCA is then used to amalgamate the individual MCA analyses of each image sequence into an overview plot for the case as a whole.

A property of the MCA and PCA analyses is that disease contours can be superimposed on the 2-D plots indicating degrees of typicality for cases of each disease [13]. A case near the centre of the disease contours is highly typical whereas cases nearer to the periphery are less typical. A further property of the plots is that the proximity of two cases of a particular disease in the plot, i.e. their perceptual proximity, indicates the similarity of the two descriptions in the original multi-dimensional space, see Fig. 1 and also Section 3.2.

In displaying cases for many diseases we adopt a largely hierarchical approach exploiting the “small worlds” metaphor [1, 9]. We divide the diseases up into small sets known as “small worlds” corresponding to sets of confusable diseases, and compute separate composite weightings for each small world. At present, subdivision of diseases into small worlds is based on the opinion of a single expert, but empirical work is in progress to verify these choices [7]. Having computed the MCA and PCA weightings for the diseases in a particular small world, we can then use the analyses to compute the separate likelihood contours for each disease in the chosen small world, see Fig. 1. The small world shown involves three broad categories of pathology. For more expert users the small world would need to be at a finer level of diagnostic discrimination.

By repeating this analysis for several small worlds, we have a set of possibilities against which a new case can be viewed. Just as a single small world can be displayed as a set of overview plots (one for image sequence and one for cases taken as a whole), so a set of small worlds can (in principle) be displayed in a composite form which presents the spatial relationship of one small world with another.

2.3 Decision Support Methodology

The following decision support methodology can be applied (with minor variations) whether the system is acting in the mode of “tutor” and offering a trainee an analysed case from the archive to diagnose, or whether the user (possibly more expert) is attempting to diagnose a case that is unknown to the system, essentially by comparing it to the others in the archive.

View MR case images. The first stage is to view the set(s) of image slices. The set is shown in the top left of Fig. 1. If case notes are available, e.g. for an image in the archive, these will not be accessible in the training system at this point, so as to reinforce the primacy of the image over other data.

Select and view small world from menu. In the current version of the system the overview plot can display only a single small world, but others will be added [7]. The rapid initial selection of hypotheses will be accommodated by the user selecting and clicking on a single button to bring the chosen small world into the overview plot.

Compare related cases from database. The overview plot is populated by dots, each dot representing a case from the archive. These dots are mouse sensitive and can be clicked on to bring the set(s) of image slices up on the screen (see the top right of Fig. 1). The user can visually compare the images for cases from the archive with the case under examination. Moreover the position of the dot in the overview plot(s) indicates through its distance from the centroid for a particular disease how typical that case is of the population for that disease.

Read clinical and presentation data. At this point it is important that the radiologist takes all the available information into account, if s/he has not already done so.

If case is no longer problematic then exit. The images and the case data may render the case in hand unambiguous to the trainee and at this point the user can exit, without further action other than to discriminate between the diseases in the selected small world if s/he can. Where the user is a trainee, and when the case under examination is known to the system, we plan to generate a reflective follow-up commentary that is sensitive to the trainee’s history of interaction with the system, the accuracy of their final diagnostic choice(s) and the process they went through in arriving at their decision.

Describe case to system. Where the case is more problematic, either because of the trainee’s lack of experience or because of its inherent difficulty, the user can engage in the more time consuming task of describing the appearance of the lesion(s) on different sequences using the menu-driven structured image description language.

See where the description lies in the small world. Using the same coefficients derived from the MCA and PCA analyses, the position of the case can be shown in the overview plot. Nearby and distant cases can then be examined by clicking on them to examine points of similarity and difference.

Check whether any other small world offers competing possibilities. It may be that the current case lies in a region of difficulty such that it is either far outside the range of typicality for any of the diseases in the small world, or in a region equidistant from two or more disease centres. In the former case, the user can investigate other small worlds to see if there are any which are both plausible. In the latter case, the ambiguous case, the system will be extended to offer advice as to which parts of the description have led to diagnostic uncertainty and/or to which further tests might be employed to reduce ambiguity.

Read off relative likelihoods from chosen small world. When the user agrees the description for the case under consideration and the best fitting small world is in view in the overview plot, then the relative diagnostic probabilities are displayed in a small bar chart.

3 Evaluation

3.1 Description Language

The analytical power of the IDL has been partly tested by its application to the differentiation of multiple sclerosis (MS) from vascular disease [4] and the effects of HIV infection on the brain [5]. Further insights into the predictive power of combinations of features will emerge as part of the continuing statistical analysis of the data, including the application of Multiple Correspondence Analysis.

3.2 Overview Plot

We have conducted a limited evaluation of the overview plot, based on the display of cases for a single disease. The evaluation [8] was carried out to investigate whether the statistically derived measures of typicality and similarity presented in the overview plot match the typicality and similarity judgements of radiologists of different degrees of expertise. This involved seventeen subjects, comprising four novices (with no knowledge of radiology), nine

intermediates (4th year medics and radiographers with some knowledge of anatomy and imaging) and four expert neuroradiologists. The average degree of agreement between human and MCA placement was in the expected order of expert (0.97), intermediate (0.95), novice (0.94). Interviews with the subjects based on a structured questionnaire indicated that they found the overview plot easy to use and acceptable as a means for retrieving cases from the image archive. The evaluation suggests that the overview plot can provide a useful teaching device, to assist a trainee in forming a mental representation of the distribution of cases of a disease comparable to that of an expert.

4 Conclusion

We have described the main components of a largely visual decision support training system derived from an existing system to teach MR image description. This is based on implementing the notion of a small world as a set of interactive overview plots, based on MCA and PCA analyses of cases from an archive. Much work is yet to be done. This includes determining how use of the system affects users' management of uncertainty, evaluating the decision-making leverage (if any) provided by the system, and evaluating the training potential of the system.

Acknowledgements

The work has been supported by ESRC Cognitive Engineering grant L127251035 and by EPSRC Medical Informatics grants GR/L53588 and GR/L94598.

References

1. C. Bishop and M. Tipping. A hierarchical latent variable model for data visualisation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3):281–293, 1998.
2. B. du Boulay, B. Teather, G. du Boulay, N. Jeffery, D. Teather, M. Sharples, and L. Cuthbert. From description to decision: Towards a decision support training system for mr radiology of the brain. In *Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making — AIMDM'99*, 1999.
3. G. du Boulay, B. Teather, D. Teather, M. Higgott, and N. Jeffery. Standard terminology for MR image description. In M. Takahashi, Y. Korogi, and I. Moseley, editors, *XV Symposium Neuroradiologicum*, pages 32–34. Springer, 1994.
4. G. du Boulay, B. Teather, D. Teather, N. Jeffery, M. Higgott, and D. Plummer. Discriminating multiple sclerosis from other diseases of similar presentation - can a formal description approach help? *Rivista di Neuroradiologia*, 20(7):37–45, 1994.
5. G. du Boulay, B. Teather, D. Teather, C. Santosh, and J. Best. Structured reporting of MRI of the head in HIV. *Neuroradiology*, 37:144, 1994.
6. M. Greenacre. *Correspondence Analysis in Practice*. Academic Press, London, 1993.
7. L. Hinkley, G. H. du Boulay, C. Goode, B. du Boulay, B. Teather, D. Teather, M. Sharples, and N. P. Jeffery. Computer based neuroradiology training using a “small worlds” approach to differential diagnosis. under review.
8. N. P. Jeffery. *Computer Assisted Tutoring in Radiology*. PhD thesis, De Montfort University, Leicester, 1997.
9. A. Kushniruk, V. Patel, and A. Marley. Small worlds and medical expertise: Implications for medical cognition and knowledge engineering. *International Journal of Medical Informatics*, 49:255–271, 1998.
10. M. Sharples, B. du Boulay, D. Teather, B. Teather, N. Jeffery, and G. du Boulay. The MR tutor: Computer-based training and professional practice. In *Proceedings of World Conference on Artificial Intelligence and Education, AIED'95*, pages 429–436, Washington, USA, 1995.
11. M. Sharples, N. Jeffery, B. du Boulay, B. Teather, D. Teather, and G. H. du Boulay. Structured computer-based training and decision support in the interpretation of neurological images. *International Journal of Medical Informatics*, 60(3):263–280, 2000.
12. 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 *Artificial Intelligence in Education: Knowledge and Media in Learning Systems, Proceedings of AI-ED97*, pages 402–409. IOS Press, Amsterdam, 1997.
13. D. Teather, B. A. Teather, N. P. Jeffery, G. H. du Boulay, B. du Boulay, and M. Sharples. Statistical support for uncertainty in radiological diagnosis. *Methods of Information in Medicine*, 39:1–6, 2000.