A Knowledge-Engineering Approach to Instructional Design

NICK RUSHBY
Centre for Staff Development in Higher Education, 2 Taviton Street, London WC1H 0BT

This paper is the report of a one-day seminar which was held to discuss the applications of knowledge engineering to instructional design and to identify areas where research is needed.

'Technological advances are like icebergs. Intriguing to observe from afar, they may become ominous as one moves closer. And, like icebergs, changes in technology are often not what they seem to be at first glance: what appears to be a large and substantial mass may on further inspection turn out to be merely a weak shell, whereas something that seems to be relatively insignificant when spotted on the horizon, may later reveal hidden qualities that make it a force to be reckoned with.'
(Stephen Kerr)

Received February 1986

1. INTRODUCTION

This paper is the report of a seminar held in London on 4 February 1986 which brought together fourteen experts in the field of artificial intelligence, education and training, to discuss the applications of knowledge engineering to instructional design. The potential of this approach has been perceived, somewhat imprecisely, for several years, but it is only in recent months that the ideas and tools have become readily available to practitioners in education and training. While the prediction of educational courseware factories,using knowledge-based tools to automate the production of computer-based learning materials, may still be some years in the future,[6] there is a feeling that we are now in a position to deliver part of the promise of technology-based learning. The two main aims of the seminar were to review the current state of the art and to develop an agreed agenda for research in the area. Within these aims lie a number of questions as follows.

(1) Which are the most significant areas in terms of their impact on education and training?

(2) What are the key research areas?

(3) Can we put timescales on the introduction of 'demonstrator' applications?

(4) What manpower resources will be needed to apply these ideas and tools?

(5) What organisational development will be needed?

(6) What are the innovation issues for education and training?

Clearly, answering these questions in full is a considerable task, which could not be completed within the compass of a single day. However, the discussions helped to clarify thoughts and the meeting ended with some consensus on future research directions.

Before reporting the content and outcomes of the seminar, it may be helpful to describe the process by which we sought to facilitate the discussions. Each of the participants was invited to prepare a short written contribution to be tabled at the start of the day. After an opening statement, setting out some of the ideas and tools that could be applied to the problem of instructional design, the discussion flowed freely until lunch time. After lunch, the participants were asked to make a list of the six research topics that they felt were most important and then divided into three groups, to discuss and collate their individual lists. The final discussion concentrated on the three resulting lists and attempted to reach a consensus on the research agenda. Readers may form their own judgement as to the extent to which the aims of the seminar were realised.

Many of the problem and ideas discussed in the seminar are equally applicable to computer-based learning and training (CBL/CBT), interactive video (IV), intelligent-computer-assisted learning (ICAL), and other applications of the new technologies to education and training. It is therefore convenient to refer to these media collectively as 'Technology Based Training' (TBT). This term has further advantages, since it encompasses training based on older technologies (such as linear video and even the overhead projector) and also serves to remind us that there is a technology *of* learning as well as the use of technology *in* education and training.

The different approaches can be categorised in terms of intervention at different levels.

(1) Intelligent Computer Managed Learning or Training (e.g. in training at Digitial Equipment Corporation,[8]

(2) Devising and validating courses taught by traditional means (e.g. at Garnet and Kingston Colleges.[2]

(3) Designing materials which are to be delivered through technology (e.g. interactive video).

(4) Intelligent CBL/CBT (e.g. SOPHIE, BUGGY; see Ref. 9.

(5) IKBS job aids changing the nature of training (e.g. DHSS regulations[7], British Nationality Act[5].

Discussions in the seminar were centred on the third and fourth of these.

2. DISCUSSION

The inadequacy of the authoring tools currently available has not inhibited the current interest in producing technology-based training materials. Despite the claims of the advocates of authoring systems, they still lack many of the facilities needed to produce 'good' instructional material efficiently and effectively. For evidence to support this assertion, we can examine the available materials to see that the majority of applications have used some elements of general-purpose programming

25

CPJ 29

languages, rather than confining themselves to authoring languages. In part this is because the authoring tools are relatively unsophisticated and because the few people who do know how to manipulate them to full effect are being inundated with work, as every major company starts to develop an interest in TBT.

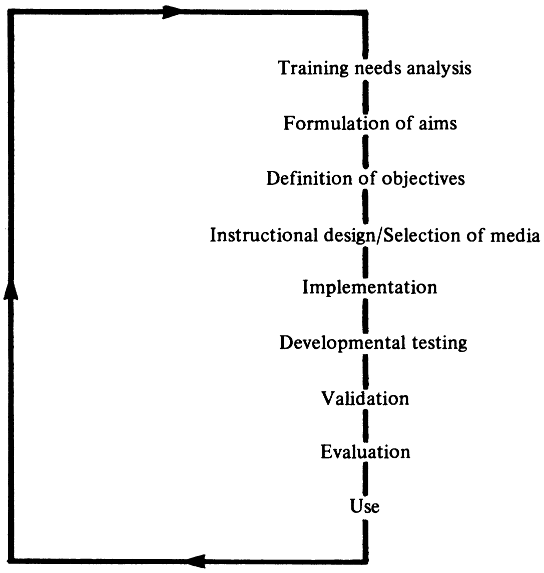

Figure 1. A systematic approach to instructional design.

Those authoring tools which are available have been aimed at the implementation and validation phases of the development process illustrated in Fig. 1. They have relatively little contribution to make to the design phase, except to provide an environment in which the designer can implement pieces of material to see what they might look like. This prototyping application is helpful, but is only one of the tools that we need.

The current generation of authoring languages, like their predecessors, are built on the principles of programmed learning. Their use is therefore limited to the simplest applications of the technology – for which they are ideal – with subject-expert trainers working in their own environments.

As users seek creative, rather than just interactive, solutions to their learning problems, the costs of design increase rapidly. All the authoring systems currently available limit the options available to the designer in some way, and the more imaginative and competent the designer becomes, the more likely he or she is to discover frustrating limitations. If the end product is unsatisfactory for the designer and fails to meet the user's ambitions, then TBT will have an uncertain long-term future. There is now a need for a new generation of authoring systems to support the new generation of creative instructional designers. These must be implemented on, and fully exploit, the abilities of the chosen delivery system. The use of expert systems is seen by some to be the next step on our road although, as some participants pointed out, you need a very liberal interpretation of 'expert systems' to see them solving problems of this kind.

3. A VISION OF A FUTURE

Let us take interactive video as an example. There is the opportunity for instructional designers to allow the learner to interact with pictures, not only after they have happened, using some academically determined questioning strategy, but whilst they are happening. Pictures are rich in information, but we have no tool to handle this wealth, or the consequent wealth of opportunities it offers to the disparate audiences who need access to it. If interactive video is to realise its potential, then we need ways of describing, in terms of the subject expertise, what pictures it needs to be most effective and, when it has the pictures, what each element of each picture means. The final outcome of the design process is the way in which all this material should be organised for the learner.

One futurist view of this process is that the system will take the subject expert(s) through a questioning process which extracts their knowledge, while allowing them considerable freedom to develop areas of particular relevance for the project in hand. It will prompt and question until it is quite certain that it has all the relevant linkages. It is quite probable that this may be an exercise undertaken by several experts with overlapping areas of expertise, and it is certain that the input will need to be validated. Once completed, the expertise area will be available and should never need to be completely re-established (unless, of course, there is a dramatic reappraisal of that area of knowledge!), although it will be updated. The complete or partial automation of the curriculum analysis leaves the designer with more time to concentrate on the learner populations and learning objectives, including: learner profiles; pre-knowledge; functional objectives; stages of learning or training; learning schedules; management objectives.

The system will begin to indicate audio visual needs (some of which may already be obvious), associating them with specific objectives. At some later date, with natural-language processing, it may also be able to handle audio and video scripts.

Such a system would result in highly interactive learning materials which could adapt to the unique needs of any student. It would continually assess the learner and, because it would automatically draw on a knowledge and experience base to supplement the structure he is following, it will seem to the student that no predetermined flow exists: the learning experience will feel personal and be personal. Although the system will be steering him, he will have the freedom to explore the knowledge base, accessing any part he feels is relevant. The system will keep track of him, measure his growing knowledge, validate its assumptions about him and the instructional strategies it is using, and report on his progress.

However, before we wave the magic wand of 'expert systems' and fly to Nirvana, there are a number of research and development issues that must be identified and tackled!

As the seminar progressed, it became possible to develop a curriculum map of the area, as shown in Fig. 2. This shows some of the possible topics and linkages: readers are encouraged to add their own as they read through this paper and further ideas occur to them.

4. DESIGN

The initial discussion focused on the process of design, and the need for the designer to have ways of describing multimedia stimuli (for example videodisc-based data). It

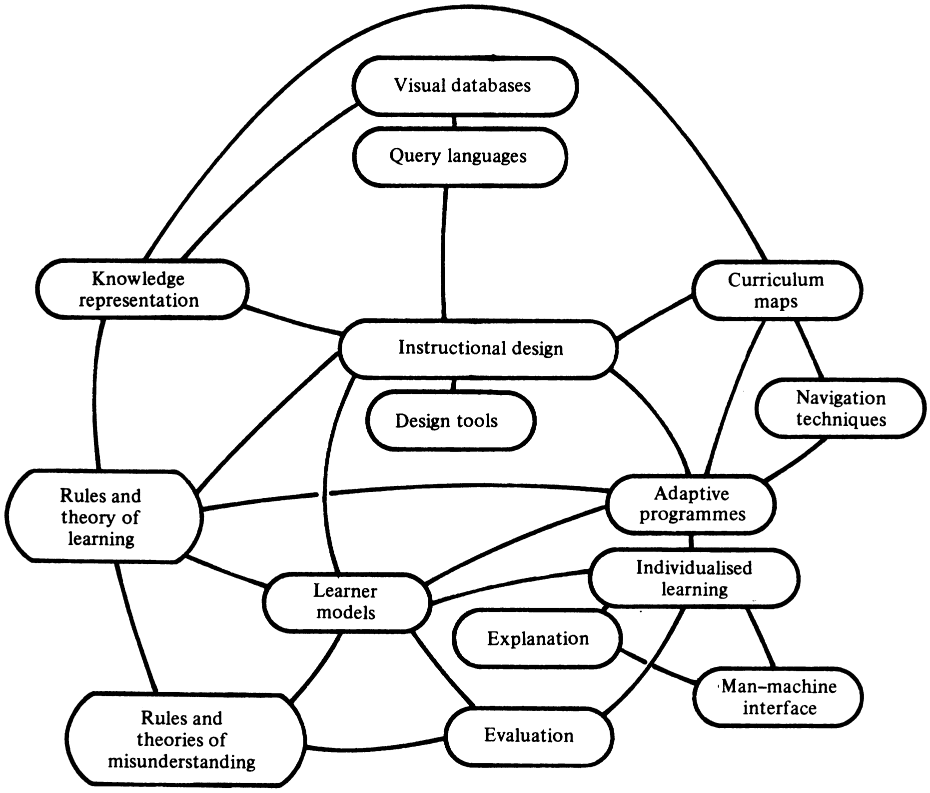

Figure 2. A curriculum map

seems likely that developments in this area will draw upon work in computer science, on relational databases and languages for querying them. Instructional design is a form of general design, and hence the principles that inform general design should also inform instructional design. Design research is particularly relevant and is concerned with:

history – why artifacts are as they are;

nature of design – its activities, organisation and apparatus;

taxonomy – classification and methods;

technology – the operation of systems;

modelling – representations of knowledge and communication;

measurement – quantitative and qualitative measures;

value – economic, moral and aesthetic worth;

philosophy and logic of discourse;

epistemology – knowing, believing and feeling;

education – principles and practice.

Designing is the purposeful transformation of an initial outline description of an artifact or system (in our case, an instructional system) into a more or less complete and detailed one by the application of information and knowledge. The transformation process is not smooth, but has discontinuities at various points. Over the continuous phases, it can be considered algorithmic and responsive to inductive and deductive reasoning. Over the discontinuities, intuition and abductive reasoning (bright, creative ideas) are essential. We are now starting to discover how knowledge-based systems can help in these processes.

In interactive video, the focus of the creative design is moving away from the subject expert and the training specialist towards the designer/producer. If evolution follows the historical precedents, then we may expect these tools to be developed for the end user – the trainer, or perhaps even the learner.

5. LEARNING

The focus then moved to the psychology of learning, and the learner's perceptions of the instruction. The learner needs to be able to navigate through the wealth of resources that TBT is able to provide, to know what questions to ask in order to access useful knowledge.

An analogy can be drawn between navigation through learning materials and across country by train and car. Traditional programmed learning materials can be compared with a train journey, where the route is predetermined by the setting of the points. Once you are on the train, then you should (eventually) arrive at your destination, without making many decisions. More sophisticated CBT is like a car journey: there are more opportunities for exploration with decisions to be made at each junction. The journey will end at one of several places, and the routes between may differ. The vision described earlier offers almost unlimited scope for exploration, with continuous decision-making by the learner and support system, and is analogous to a four-wheel-drive vehicle on a cross-country expedition – it is as beneficial (for learning) to travel inquisitively as it is to arrive.

Clearly, there is a need for a better understanding of the learning process, and a need for a language with which to describe it. Although considerable work has been done, there is some difficulty in applying the results to the kind of systems that are now proposed. Instructional design is like other areas of education in being deficient in testable theory, in powerful tools and in powerful hypotheses. This has led to over-reliance on the cycle of test–revise–retest (discussed and found wanting above); to a lack of questioning of the value and validity of material provided by teachers; and to its corollary, a teacher-dominated curriculum. In addition we are poorly placed to diagnose errors. Any learner errors which are found during tests are often viewed as random, or merely indicative of the need to change the teaching material: they may, in fact, be consistent applications of imperfectly understood procedures. Most of the current TBT materials concentrate on success in mastery learning and do not linger over errors. We can learn much from errors and we should pay them more heed.

Such problems may be reduced by using ideas and techniques from machine learning and knowledge engineering. Those ideas and techniques promise improved understanding:

– how to elicit and ensure the consistency of expertise (including subject-matter expertise and teaching expertise);

– the effect of learners' prior knowledge and expectations on their interpretation of advice and information;

– how to communicate learners' viewpoints and abilities to teachers;

– how new knowledge (e.g. concepts or steps to accomplish a task) is acquired, classified, integrated and operationalised;

– how and when existing knowledge is revised, classified, extended, rejected (e.g. to resolve conflicts, integrated and operationalised;

– the relative efficiency of different ways of learning;

– the effect on learning of the quality of new information (e.g. the effect on an inductive learning program of a training set which includes irrelevant examples or randomly ordered examples);

– the effect on learning of the quantity of information (e.g. how many examples to provide);

– the effect on learning of the method in which information is presented;

– how learning difficulties can be anticipated, diagnosed and/or remedied;

– whether learning can be assessed more objectively;

– the nature of analogy and how analogies are recognised.

6. AN AGREED RESEARCH AGENDA

The first half of the seminar identified a number of broad areas of concern, as shown in Fig. 2. During the afternoon session these were shaped into the form of an agenda by a three-stage process, whereby fourteen individual lists were amalgamated into three and then into one agreed list. There was a remarkable degree of unanimity in the lists, which we may interpret as an indication of the urgency of the research. The final, agreed list presented here is not in any specific order: some areas may require more time and resources while others may be resolved more easily; some should be accorded a high priority; but all are important. Readers are invited to comment on the items and to suggest priorities.

(1) Learner models: the identification and refinement of helpful models of learning and learners, of strategies and patterns of errors, together with a taxonomy and a language for talking about them. As we have seen earlier, most TBT follows a programmed learning-like style which is not necessarily effective for all learners and all topics. As Lewin has remarked, 'Nothing is as practical as a good theory.' TBT needs a broader repertoire of learner models on which to base the new generation of learning materials.

(2) Learning structure: the development of ways to assist and encourage the learner to access and negotiate pathways through knowledge which may have a greater or lesser degree of structure; techniques and languages for signposting and for navigation.

(3) Explanatory systems and evaluation systems. There are strong links between the processes of explanation to the learner and assessing whether he or she has acquired some particular piece of knowledge. The process of evaluation is, of course, somewhat recursive, since it not only examines the learner, but also the instructional materials and the evaluation instruments themselves.

(4) Knowledge elicitation, elucidation and representation. The provision of better tools for the designer and author, to help in the process of knowledge engineering.

(5) Curriculum map codification. One way of representing a knowledge area is in the form of a map (as, for example, in Fig. 2). How can these be coded so that they can be used within automated instructional systems?

(6) Complex database representation: rules and heuristics for handling and querying very large collections of media-sensitive information. Initially developed for the designer, these could (should) also be made available to the learner.

(7) Modular (generic) design elements and production tools. The automated courseware factories envisaged for the future, and producing generic and bespoke learning materials, will require rules for creative visualisation and a range of tools which apply much of the research identified above.

(8) Man–machine interaction; its effectiveness and guidance for users. The results of earlier research in this area must be applied to technology-based learning to see whether the theories need to be modified. Are there significant differences between well-researched interactions for, say, computer-aided design in engineering, and the interaction between a learner and his or her technology-based resources? To what extent can the user be guided by an expert system? Are different forms of interface more or less effective?

7. WHERE DO WE GO FROM HERE?

The application of knowledge engineering to the various aspects of instructional design has been identified as a key area in educational and training technology. It has been growing in importance over the past months and is likely to dominate research in machine-mediated learning in the near future. Two important outcomes of the seminar were the identification of an agenda for research and development in the area and an agreement to improve the dissemination of information about the work currently in progress. Experience has taught us that providing an effective information service requires significant resources, for it is difficult to set the boundaries of the subject area and of its users. The scope of the service grows rapidly and, as the demand on the resources increases, it becomes necessary to recover the costs from the users. As a short-term measure the author has agreed to act as a clearing house for information, including the publication of a directory of the projects represented at the seminar. In the longer term, it may be necessary to seek additional funding.

The application of knowledge engineering to learning is an interdisciplinary subject, drawing heavily on both computer science and education technology. The research and development which is needed must involve both communities, and so the way forward must lie in collaboration. We can identify at least three communities working in this field in the UK with relatively little communication between them: those working under the Alvey banner, those working in educational and training technology, and those whose primary interest is in computing. The problem is to draw all three communities closer together so that there is more opportunity to share experiences and information.

The identification of a research agenda is only the first step towards securing the resources needed and getting the work done. The agenda discussed earlier is to form the basis for a workshop at the forthcoming Educational Technology International Conference (ETIC'86) to be held at Heriot Watt University, Edinburgh in April, and the theme will be taken up again a year later at ETIC'87 in Southampton. In the meantime, follow-up seminars are being planned.

Acknowledgements

Thanks are due to the British Computer Society for sponsoring this seminar, to the Editor of its *Journal*, Peter Hammersley and to John Duke for organising the event. I have deliberately avoided attributing specific comments to individual participants since to do so would break up the flow of the report and, to a certain extent, belittle the synergy that was created and which gave rise to some of the ideas. My thanks go to the participants – Jonathan Briggs, John Coffey, Jim Cunningham, Richard Ennals, John Lansdowne, Paul Lefrere, Bill Plummer, Keith Roach, Peter Roberts, John Self and Roger Wilson – whose names, by rights, should appear as co-authors. The responsibility for my perceptions – and possible misunderstandings – of what they said lies, of course, with myself.

Biographical note

Nick Rushby has been involved with the use of computers in education and training for over fourteen years. From 1978 to 1984 he directed computer-based learning activities at Imperial College, running CEDAR, the national information service on the use of computers in education and training, and CITAR, a unit providing workshops and consultancy on the use of new technologies in training. In the summer of 1984 he moved to the London Centre for Staff Development in Higher Education, where he is responsible for new technologies in training. He is author and editor of a number of books and papers on CBL and CBT, and is an active member of the Association for Educational and Training Technology.

BIBLIOGRAPHY

Readers who wish to pursue this topic further may find some of the following books, reports and papers helpful.

1. P. Barker and H. Yeates, *Introducing Computer Assisted Learning*. Prentice-Hall, London (1985).
2. B. Cameron, Computer aided educational consultancy. *MSc Dissertation*, Imperial College Department of Computing, London (1983).
3. CERCLE, Centre for Research on Computers and Learning: Project descriptions. University of Lancaster, Bailrigg, Lancaster LA1 4YR, UK (1985).
4. Convergent, *Intelligent interface for integrated knowledge systems: management summary*. Convergent Communications Limited, 26 Eccleston Square, London SW1V 1NS (1985).
5. H. T. Cory, P. Hammond, R. A. Kowalski, F. Kriwaczek, F. Sadri and M. Sergot, *The British Nationality Act as a Logic Program*. Imperial College Department of Computing, London (1984).
6. J. F. Duke, *Educational Environments for the 1990s*. Unpublished report to the Council for Educational Technology, London, p. 73 (1983).
7. P. Hammond, *Representation of DHSS Regulations as a Logic Program*. DoC Report 82/26, Imperial College Department of Computing, London (and in Proc. BCS Expert Systems Conference 1983) (1982).
8. J. Heines, Basic concepts in knowledge-based systems. *Machine Mediated Learning* **1** (1), 65–95 (1983).
9. P. Lefrere, R. H. W. Waller and P. Whalley, Expert systems in educational technology. In *Aspects of Educational Technology*, vol. xiv *Educational Technology to the Year* 2000 edited R. Winterburn and L. Evans. Kogan Page, London (1980).
10. R. McAleese, Some problems of knowledge representation in an authoring environment. *PLET* **22**, (4), 299–306 (1985).
11. T. O'Shea and J. A. Self, *Learning and Teaching with Computers: Artificial Intelligence in Education*. Harvester Press, Brighton (1983).