DOCUMENT RESUME

ED 317 157 IR 014 030

AUTHOR	Cumming, Geoff
TITLE	Artificial Intelligence Applications to Learning and Training. Occasional Paper--InTER/2/88.
INSTITUTION	Lancaster Univ. (England). Dept. of Psychology.
SPONS AGENCY	Economic and Social Research Council, Lancaster (England).
PUB DATE	Aug 88
NOTE	32p.; Information Technology in Education Research Programme.
PUB TYPE	Collected Works - Conference Proceedings (021) -- Reports - Descriptive (141)
EDRS PRICE	MF01/PC02 Plus Postage.
DESCRIPTORS	*Artificial Intelligence; *Computer Assisted Instruction; Foreign Countries; Higher Education; Media Research; *Seminars; *Training
IDENTIFIERS	*Research Trends; United Kingdom

ABSTRACT

This report summarizes and interprets the discussions at a seminar on artificial intelligence (AI) training domains and knowledge representations which was sponsored by the United Kingdom Training Commission. The following broad areas are addressed: (1) the context, process, and diversity of requirements of training and training needs; (2) defining AI, expert systems, and prospects in AI; (3) the origins, recent approaches, and current research directions in the use of computers to enhance learning; and (4) AI applications in training. Implications for the Training Commission are then considered. A 9-item annotated bibliography is included, and an update of the Training Commission's program and a list of seminar participants are appended. Other publications of the Economic and Social Research Council relating to information technology and education are listed, along with ordering information. (MES)

/E/S/R/C

Economic and Social
Research Council

TRAINING
Commission

ARTIFICIAL INTELLIGENCE APPLICATIONS TO LEARNING AND TRAINING

Occasional Paper -- InTER/2/88
August 1988

Edited by:

Geoff Cumming

Department of Computing, Lancaster University

Department of Psychology, LaTrobe University

Information Technology
in Education Research
Programme
Department of Psychology, University of Lancaster, Lal 4YE

AI Applications to Learning and Training

PREFACE

By the summer of 1988, the Artificial Intelligence Applications to Learning Programme of the Training Commission (formerly the Manpower Services Commission) had been in existence for a little over a year. The first round of projects and other activities had been funded, some were nearing completion.

In addition to searching for strategies of timing and approach for a campaign to increase the awareness of AI potential amongst the training community, the External Advisory Group for the Programme were seeking the most effective strategy for the next phase of development projects.

It was agreed to hold a small invited seminar of trainers, cognitive scientists and educationalists to explore the possibility of forming a 'map' of training domains and to associate knowledge representation formalisms (ways of representing knowledge in a computer system) with those domains. It might then be possible to view the existing projects in terms of such a map, to identify the most promising domains and use this perspective to help place priorities on future projects.

This ambitious task was not achieved. However, in the attempts to move partway in this direction, the discussion did help to clarify a number of issues and provide the Training Commission with some guidelines for future action. This paper aims to describe these and help in promoting an understanding of key aspects of the nature of the potential of AI applications to training.

We are grateful to those who participated in the seminar and particularly to Dr. Geoff Cumming for preparing this paper as a report on what was said and on what might have been said.

Professor R. Lewis,
ESRC-InTER Programme

Dr. John Gillingham,
Learning Technology Unit,
Training Commission

CONTENTS

EXECUTIVE SUMMARY

The initial aim of the seminar was to consider how AI can be applied to training by making a taxonomy of training domains, then associating with each an appropriate knowledge representation formalism *. This ambitious and conceptual exercise remains for the future. In the event the seminar discussed a variety of perspectives on how AI can be used to enhance learning. This report is an interpretation and extension of those discussions.

Training perspective

Training needs are changing, and new technologies should be exploited to enhance training. Training is needs driven and output oriented: new methods will be adopted only if they are proved effective in practice. The classical systems model of training is a specify-plan-prepare-deliver-evaluate-feedback cycle. New methods could contribute at any point in the cycle.

Training includes a great diversity of goals and methods, going well beyond the acquisition merely of the skill needed to carry out a single task. Much training can be regarded as promoting expertise, including higher-level strategic and judgemental skills; this richer view of training should result in employees with more flexible and adaptable abilities.

AI perspective

Artificial intelligence (AI) embraces expert systems, but much else as well. There has been over-selling and disillusionment. AI research is making only slow progress, but tools and techniques based on AI ideas are proving to be useful and commercially successful in many selected applications.

AI applications to learning

Research on intelligent educational systems (IESs) has explored a wide variety of ways that AI tools and techniques might be used to support learning. The aim is not to emulate the full performance of a skilled human tutor, but to enhance a selected part of the teaching and learning interaction between learner and computer.

* Such a formalism is essential to any AI application. It is the way in which human knowledge is expressed in a formal manner so that it can be processed by a machine. There are a wide variety of such formalisms.

Important ideas and approaches include:

- representation of the knowledge in the domain of interest separately from teaching expertise and knowledge about the learner;
- use of knowledge of each learner's progress and misconceptions in order to individualise the IES response;
- provision of some (limited) natural language understanding capability, to improve communication with the learner;
- giving the learner a 'learning environment' or set of computer-based modelling tools, so that he or she can 'learn by doing';
- using powerful workstations with advanced graphics capabilities to present a simulation of the target machine or system to the learner;
- designing simulations, or other IESs, so that the learner can build a good conceptualisation, or 'mental model', of the target machine or knowledge;
- addition of an intelligent advisor or coach to a simulation or learning environment, so that the learner's exploration can be guided, and his or her questions can be answered;
- addition of 'intelligent help' facilities to software packages or computer-based machines, so the user can learn on the job;
- judicious use of an expert system to enhance the preparation or delivery of conventional training;
- using the apprenticeship metaphor, in which the learner does the task, with guidance, assistance and evaluation;
- use of AI tools to allow rapid prototyping of computer-based training (CBT).

Conclusions

There is great potential for AI applications to enhance training and learning. IES research gives many models that can guide applications. but application of the results of this research has not been rapid, partly because there is a shortage of expertise in IES development.

Training Commission activities aimed to raise awareness of the training applications of AI, and to illustrate such applications by financing projects to develop demonstration software, is well justified. Care must be taken to base campaigns to increase awareness mainly on working demonstration software that can be seen to be useful. Demonstrators need to be carefully selected and planned, involve industrial or other appropriate collaboration, and address a practical learning need. It is essential that demonstrators be informed by IES research; the best way to achieve this, given the shortage of IES expertise, may be to engage leading IES researchers as consultants.

1. INTRODUCTION

The seminar 'Training Domains and Knowledge Representation' was sponsored by the Training Commission (TC) and organised by the Information Technology in Education Research Programme of the Economic and Social Research Council, in the context of the TC's Programme on Artificial Intelligence Applications to Learning.

The TC's Programme on Artificial Intelligence Applications to Learning represents a commitment of some £3.2M over 1987-1990. Any activities or projects arising out of discussions at the seminar would need to be completed by March 1990 if they are to be accomplished within the current TC Programme. The current state of this Programme is outlined in Appendix A.

Participants were training managers from commerce and industry, computer people with experience in artificial intelligence (AI), academics with interests in computing and learning, and members of the Learning Technology Unit of the TC. A list of participants is given in Appendix B.

Most discussion consisted of presentations of a wide variety of perspectives on training, and on the possibilities for using AI tools and techniques to enhance learning and training. Participants described, on the one hand, what industrial training needs are, and on the other what AI and other computer techniques can offer. One interesting outcome was the realisation of how varied the perspectives are. This conclusion emphasises the need for further efforts to improve communication, especially between industrial trainers, and researchers who are studying AI applications in learning.

The TC's Programme is concerned with Higher and Further Education, as well as training. The emphasis in this report on training reflects discussion at the seminar, but most of the general issues and conclusions have applicability also more widely in education.

As indicated by the title of the seminar, the original intention was to try to identify training domains and to associate with each a suitable knowledge representation formalism. Such a mapping would emphasise the breadth and diversity of training needs, and also could guide the specification of demonstrators and other projects within the Programme. In the event the seminar did not make a concerted attack on this challenging problem, which remains for future exploration. The report does, however, include some brief discussion of the issue.

The seminar had no formal agenda and no decisions to take, and it spelled out no specific conclusions or recommendations. This report is based on the discussions recorded over the two days of the seminar, but is necessarily a personal interpretation, and includes elaboration of some issues mentioned only briefly at the seminar.

Following an outline of the seminar's aims, there are sections describing in turn perspectives on training needs, on AI, and on AI applications for learning. Then more specific possibilities for AI applications in learning are suggested, and conclusions drawn in relation to the TC Programme.

Discussions at the seminar were based on five implicit goals:

- to discuss training needs, and the possibilities for AI applications, so as to increase the understanding by participants of perspectives other than their own;
- to consider what knowledge representations, or other AI techniques, might be appropriate for various identified training domains, and so to assist the TC in seeing where current projects in the AI initiative were placed and what new areas may be promising for future demonstrator projects;
- to consider in general what applications of AI in training seemed feasible and promising in the short and medium term;
- to suggest what steps the TC might take to increase awareness of how AI can enhance training; and
- to raise any further points considered relevant to the success of the Programme.

2. TRAINING AND TRAINING NEEDS: A PERSPECTIVE

2.1 THE CONTEXT

There is widespread realisation that rapid technological, industrial, social and demographic change is leading to increasing need for industrial and commercial training. The training needed, and the means for providing this training are also changing. At the same time training departments are feeling pressures to increase efficiency and effectiveness. There is, overall, insufficient training, with many organisations not providing any systematic training.

New technologies should be able to enhance training whether by increasing effectiveness or reducing costs. Some use of new technologies is occurring, for example, linear video, videodisk and computer-based training (CBT), but there is some disillusionment and uncertainty as to what promises will actually be fulfilled.

Specific influences on training include:

- the shrinking population of young people;
- the prospect of Europe 1992;
- increasing need for retraining, and the concept of continuing lifelong education;
- the needs of special groups e.g. women returning to work, the unemployed;
- the need in some cases that training be designed, prepared and delivered very rapidly;
- the need for workers to be more flexible, so training should not be narrow;
- severe and apparently chronic shortages of skilled people in some areas, especially in the new technologies themselves;
- the need to train and retrain sufficient trainers, especially in new technologies where there are shortages anyway.

Much training at present, including CBT, is seen as somewhat ad hoc, limited and uninteresting. It would be a great advance if training could be more adaptive to the needs of the individual learner, without the prohibitive cost of personal human tutors. The goals of AI, and claims made by some for it, include this sort of individual adaptability, so surely AI applications in training hold great potential?

Above all, training must be recognised as a practical applied activity, carried out under pressure. Learning outcomes and later job performance are what is looked for; teaching and training are means to these ends. The time scale is short, the pressures are day-to-day. Any new techniques or tools must be able to deliver, and to be seen to deliver. Claims and promises do not suffice: to be taken up, a new method must pass the 'so what' test. We must avoid having a claimed solution looking for a problem.

AI applications are likely to be considered as an extension of other new possibilities, such as interactive video, videodisk, satellite communications, and CD-ROM.

Any awareness campaign, whether to senior managers or to training managers themselves, must be based on working systems and real applications. An industry is likely to be persuaded best by an application that can be seen to apply to that industry. Without these it would be better to hold back for the present. The promise is there but the time may not yet be ripe; the demonstrators must come first.

2.2 THE PROCESS

Training should be considered as a full process, typically with the following steps:

1. Identify the training need.
2. Analyse the training need.
3. Specify what training is to be provided.
4. Write detailed training objectives, including specification of abilities that may be assumed in the trainees.
5. Design the training.
6. Prepare materials, train trainers.
7. Administer training.
8. As a result of experience and evaluation at 7., loop back to steps 4., 5., or 6. Repeat, until training is sufficiently refined and successful.
9. Carry out external validation and evaluation; loop back earlier if required; note improvements to be incorporated next time training is to be designed and prepared.

It is important to note the full range of activities involved in planning, preparing, delivering and evaluating training; there may be scope for AI to contribute at any point in the process. A detailed analysis of the training process, and the job of a trainer may therefore be required.

While the scheme described above may be the ideal, it does not describe universal reality. Ways to make the ideal easier to achieve could be valuable.

Good authors of training materials can produce fairly effective courses. using CBT or other modes. But production of CBT materials. even with current authoring tools, can be a tedious, repetitive, expensive task. New methods are needed.

Also, best practice is not universal practice, and much training material, including much CBT, is unimaginative, inflexible and dreary.

The great strength of a human trainer is flexibility: if one approach does not work, the trainer will try an alternative approach. New technology-based methods having this capability would be very valuable.

2.3 *THE DIVERSITY OF REQUIREMENTS*

No single view of the nature of training should be allowed to restrict our thinking. There is a great diversity of training needs, going well beyond, for example, the learning of specific procedures needed to operate a perform a specific task. It needs to be recognised that human expertise is complex and deep, and includes higher level skills of judgement, problem solving and being able to take initiative. Training for the various levels of management, for example, requires attention to a wide range of abilities. It is also worth noting that although the TC has 'training' in its title, the name of its programme on AI applications refers more broadly to 'learning'.

Two further examples illustrate the diversity of problems that must be addressed by trainers:

Steel-making

British Steel is undertaking a large exercise in setting competency standards for every class of job, at all levels. This is a complex task, especially as the managers of the 40 plants have differing views as to what skills are needed for particular jobs. Traditions, equipment characteristics, and job organisation also differ from plant to plant. Setting competency standards may be regarded as an early step in the specification of training needs. Might AI be able to help?

Health & safety

The new Health & Safety legislation incorporates a vast mass of tedious, but precise and necessary detail. A mere statement of this information, or even retrieval of the parts of it relevant in a particular case, is not sufficient. There is, rather, a practical need for interpretation and conceptualisation, and judgement as to what information is applicable in some situation. These considerations apply both to the application of the legislation, and to training. Could some sort of expert system help?

2.4 *THE TRAINING PERSPECTIVE: A SUMMARY*

- At present there is insufficient training, and the training that is offered is often poor. Demographic, social, political and technological developments are all changing training requirements and increasing the need for training.
- There is, therefore, need and scope for improved methods in the preparation and delivery of training. Training is, however, needs driven and output oriented: new methods will be adopted only if they are proved effective in practice.
- The classical systems model of training is a specify-plan-prepare-deliver-evaluate cycle. New methods could contribute at any point in the cycle.
- Training is much more than acquisition of the skill to operate one machine. The great diversity of training needs, and of types of training activity must be recognised. There are, therefore, many opportunities for new techniques to be used to improve training.

3. ARTIFICIAL INTELLIGENCE: A PERSPECTIVE

3.1 WHAT IS AI?

Minsky's early definition of AI is still widely accepted. It states that: 'Artificial intelligence is the science of making machines do things that, if done by humans, would be regarded as intelligent'. Important aspects are the ability of a system to adapt, to learn, to use knowledge in novel ways, and to cope with unforeseen circumstances.

The definition above is in terms of the behaviour of the whole system, but it is also important to consider how the system works. A chess program, for example, may play quite impressively but be based on 'brute force' methods i.e. simply be able to consider a vast number of possible future moves very quickly. Such a program would not be classed as having much intelligence. A different program may rely on analysis of the board position and judgement about good moves that is much closer to the way human experts play chess. This program may not perform as impressively as the first, but may be regarded as having some intelligence because its methods give scope for learning and improvement.

AI as a research field is an active, argumentative and changing enterprise. Debate continues as to what is, and what is not AI, and what definition should be used to decide. In addition, changes over time must be recognised: things that were regarded as AI a few years ago may now be taken for granted as routine and non-intelligent.

An example is the range of tools, such as object-oriented programming, and graphics-based workstations, that were originally developed as part of AI research and which embody AI ideas to some extent, but which are now becoming quite widely used in routine software development.

Concern about definitions of AI, and connotations of using the label 'AI', led the TC (then the MSC) to debate whether its programme should be renamed. But 'knowledge-based', or any other single label, did not seem clearly better, quite apart from the visibility problems of a name change.

In practice, 'AI' is used to refer variously to desired characteristics of systems - adaptability, for example - but also to techniques, such as explicit knowledge representation, use of powerful tools, and use of heuristics. In the context of the TC's Programme, it is probably best to take a pragmatic attitude to what is meant by 'AI', while remaining extremely wary of any claims or assumptions that rest merely on the use of a glamorous label: not just anything may legitimately be called AI.

3.2 EXPERT SYSTEMS

Expert systems are the most widely known examples of AI applications; for some people the two labels are close to synonymous. It is important to note the limitations of expert systems, and also to realise that expert systems comprise only one small area within the great range of topic areas that is AI. Understanding of natural language, knowledge representation, machine learning, robotics, and cognitive modelling are examples of other areas in AI with relevance for training.

An expert system is a computer program containing a body of topic area knowledge, some ability to reason with that knowledge, and an interface that allows a user to consult the knowledge and seek answers to problems in the topic area. Expert systems are typically built by eliciting knowledge from a human expert and expressing this in the program. The program can contain judgemental and relational and uncertain information, as well as clearcut facts.

But expert systems have at best only the glimmerings of human expertise. The large expert systems that have received the press coverage (MYCIN, PROSPECTOR, DENDRAL) have taken enormous effort to develop, yet can show just some aspects of expertise and that only within a narrow field.

Many systems are now built by using a 'shell' - a program providing logical reasoning, knowledge formalism and interface facilities to which domain knowledge can be added to form an expert system. Using a shell makes it quicker to build a system, but the result is extremely limited. Many knowledge domains of practical interest can be only superficially represented in the simplistic rule-based formalisms typically available.

Some small expert systems, whether developed from scratch or by using a shell, are proving to be effective in use, and are commercially successful. But their degree of expertness is tiny. Usefulness need not imply expertise or intelligence. Developing expert systems that can give good explanations of their answers and reasoning remains an active and challenging research enterprise. The use of the word 'expert' may be unfortunate, and must not be allowed to undermine our normal critical approach.

It is vital to realise that an expert system can be used in a wide variety of ways. The value may come from building it, exploring its knowledge, or using it with imaginary problems. It should be looked on as a resource, or tool, just as we look on books or films or simulations as - very different - resources or tools. An important contribution may be the discipline the expert system imposes as we try to elicit the knowledge from a human expert, and formalise it in the computer system. Expert systems should be seen as tools to assist in tasks normally carried out by an expert. They should not be seen as replacing the need for human involvement or decision-making.

3.3 THE PROSPECTS FOR AI

A distinction should be made between fundamental research in AI, and applications. Many academic research workers in AI would say that, although research progress continues, large and fundamental leaps forward are not being made and are not imminent. Better software tools are being developed, as well as better and cheaper hardware, and established techniques are being used to build systems with impressive performance. But the main advance in our understanding is an increased respect for the complexity and subtlety of human intelligence, and a realisation that the AI road is long and hard.

The prospects for AI applications have of course been the subject of enormous and widely-known hype, over-selling, claims and counter-claims, and disillusionment. The prosaic reality is that applications for some years to come will be of tools and techniques that already exist. Many useful applications will be made, and software engineering practice will rely increasingly on hardware, software and interfaces that have AI parentage. Systems will become easier to use, especially for untrained people, and some systems will be easier to de[illegible] but progress will be evolutionar,. AI hype must not be allowed to over-rule our caution and critical attitude.

3.4 *THE AI PERSPECTIVE: A SUMMARY*

- The AI enterprise is an endeavour to build systems that show the flexibility, learning, and other characteristics of human intelligence. Progress has been limited and slow, but powerful tools and techniques for knowledge representation, software development and interaction with users have been developed.
- In expert systems the knowledge base and the reasoning module are separated. Shells allow reasoning and interface modules to be reused with different knowledge bases. Expert systems can be practically useful in some limited applications, although their ability to explain their answers is very limited.
- AI comprises much besides expert systems. Natural language understanding, robotics, and cognitive modelling, among others, are areas with relevance for training.
- Over-selling must be resisted, but there are now tools and techniques available that have many useful applications - in training and elsewhere - even if the resulting systems have only glimmerings of 'intelligence' or 'expertness'.

4. COMPUTERS TO ENHANCE LEARNING: AN AI PERSPECTIVE

4.1 *THE ORIGINS*

In the 1950s Skinner proposed teaching machines, based on behaviourist principles, that gave the learner a linear sequence of items. Crowder introduced branching programs. Gordon Pask built the SAKI typing tutor, which adapted to the individual learner by choosing the letters giving the most trouble and presenting them for practice. Touch-sensitive screens were introduced.

In the 1960s, Smallwood's geometry program was the first self-improving tutor: it conducted experiments on its learners, trying to find better ways to route a learner through the material. Uttal, Uhr and Wexler demonstrated how stored algorithms could be used to generate tasks for the learner, thus avoiding the necessity to enter every item in advance. Papert and Feurzeig developed from *LISP* the computer language *LOGO* for educational use.

In the 1970s, Hartley and Suppes introduced task analysis, and the notion that task difficulty could be matched to the learner's needs. Alan Kay specified Smalltalk, the first object-oriented programming language. Kimball made the first use of an expert system as part of a teaching system: his tutor for symbolic integration learned from its learners by storing any learner's solution that was better than any solution known to the system. Self, in 1974, spelled out the value of incorporating an explicit model of the learner, so that the system could be adaptive to the needs of that individual.

The electronic trouble-shooting program SOPHIE was based on simulations of electronic circuits. It had an impressive natural language capability, that is it could converse with the learner in something approaching normal English.

In parallel and largely independent of this work, a great deal of research and development has been undertaken on 'non-intelligent' uses of computers to enhance learning through simulations and other exploratory software mostly in schools and colleges. This important work is not discussed in this brief history of instructional and tutorial applications though the experience of such work is now being recognised in the increasing use of simulations in training.

4.2 RECENT APPROACHES

'Intelligent Computer Assisted Instruction' and 'Intelligent Tutoring System' (ITS) are the terms most widely used for systems using AI ideas or techniques. ITS, in particular, is an unfortunate term because it suggests a limited one-way tutoring style in which the system teaches the learner. It may thus have the damaging effect of narrowing our conceptions of how AI applications can support learning. The broader term 'Intelligent Educational System' (IES) is preferable and will be used here.

The traditional way to describe the structure of an IES is to identify three major components:

- domain knowledge i.e. knowledge in the subject area or of the task chosen for learning;
- teaching knowledge;
- a learner model, so that response can be adapted to the current needs of the individual learner.

It seems logically necessary that any reasonable IES should contain, at least implicitly, knowledge of these three types.

Expert systems

The pioneering expert system MYCIN was based on several hundred production rules; a production rules has the form 'if (condition) then (action)'. Clancey started with MYCIN, then developed GUIDON, NEOMYCIN and other systems to explore the use of expert systems to support learning. The major conclusion was that, although MYCIN performed creditably at medical diagnosis of infectious diseases, it was of only limited use in a tutoring system. Learning requires explanation, and an expert system based on a simple set of if-then rules can only give explanation of limited scope, possibly understandable by experts but not by novices to the subject. Much richer knowledge representation, including broad conceptual knowledge and problem solving skills and strategies, is required for a good tutoring system.

Simulations

Two large computer-based simulations intended for training use are STEAMER, and the RECOVERY BOILER TUTOR (RBT). In each case there are impressive facilities for displaying the state and behaviour of the system. The user can select from many schematic diagrams showing the workings of aspects of subsystems or the whole system. Icons (small pictures) and displays are used to represent salient aspects (levels, voltages, pressures, etc) at many points. The user can intervene, and observe how the system reacts to the change made.

On top of the simulation is an advisor or coach (well-developed for the RBT; in development for STEAMER) which can guide, advise and tutor a user who wants to learn about the system. A variety of interactions with the learner is supported, for example:

- the coach sets up a problem condition, the learner tries to manage and correct the problem, while receiving advice from the coach;
- the learner can take the initiative and try any change, observe what happens, and hear commentary from the coach;
- the coach can take the initiative and guide the learner through the sequence of actions that an expert might take to correct a problem;
- the coach can introduce a fault for the learner to diagnose, if necessary calling for help from the coach;
- a pair of learners can collaborate in using the system.

The wide variety of types of interaction is a very important feature of these systems. Another notable aspect is the interface. Powerful workstations with mouse, large graphics screens and windowing software were used. The graphical representation of system functioning was designed with great care. A learner can interact very closely with the simulation, and see immediate links between changes introduced and effects on system behaviour.

Adding a coach to a simulation requires not only the representation of extra knowledge about the system and how it works, but also some knowledge about advising and teaching, and the ability to build and maintain a model of the individual learner.

These systems are impressive, and are reported to have been used successfully to give training in realistic settings. They were large-scale development projects; an important question is whether future systems can be built any more quickly with less cost. Improved tools will help, but there seems no way of escaping the need to draw on the best human expertise in the area, and a great deal of technical knowledge about the target system being simulated.

Design of simulations - and other IESs - is based on the notion that the goal of learning should be deeper understanding by the learner. In most cases we want to go beyond mere skills training to have the learner develop a good 'mental model' of the target system and how it works. Such an approach to training should give more flexible abilities, and lead to staff being more adaptable, and better able to cope with unforeseen situations and with new versions of the equipment.

The screen displays of the simulations described above, for example, are not necessarily designed to give close physical fidelity with the target system, but to have good conceptual fidelity with the mental model, or conceptualisation that we would like the learner to develop.

Anderson's tutors
Anderson's *LISP* tutor is based on a large set of production rules. The rules model correct *LISP* programming skill, and also many errors made by learners learning *LISP*. A learner works through a series of programming exercises. The system monitors the learner's performance and intervenes when an error is made. It also answers questions and gives advice if asked. The *LISP* tutor is reportedly successful in tutoring beginners, and is now available commercially. It can be criticised for the very tight rein it keeps on a learner, which gives an oppressive tutoring style. It generally recognises only one correct way to solve a programming exercise, so a learner has no scope to explore alternatives, or to make a sequence of errors then try to recover. Anderson is also developing similar tutors for high school geometry and algebra.

Learning environments
Papert advocated an approach entirely different from that of CAI and intelligent tutors. Recalling very early notions of liberal education, he insisted that the learner should have a set of tools, or a domain for computational play, and be encouraged to explore, experiment, discover things, and, most fundamentally, construct his or her own understanding, rather than be the passive recipient of someone else's understanding.

'The child should instruct the computer, rather than the other way round' was one of Papert's slogans. In practice this meant that Papert wanted the learner to use *LOGO* for simple programming projects; the discipline of expressing ideas in the *LOGO* formalism, and debugging *LOGO* programs so that they behaved as intended, would promote, Papert claimed, important higher-level abilities, such as problem solving, project management, and the ability to use errors to guide improvements.

Attempts to put Papert's ideas into practice have had mixed success, and evaluations of the use of *LOGO* have also been mixed, but the notion that learners should be given the tools and the scope to take initiative and build things themselves has been influential.

4.3 CURRENT RESEARCH DIRECTIONS
The projects and approaches described above can be regarded as steps on the path to the fully powerful IES. This system would include multiple, rich representations of the domain knowledge so that it could give good explanations at a number of levels; it would contain teaching knowledge to support a variety of styles of interactions with the learner; and a learner modelling component that could match response to the current understanding, misconceptions and preferences of the learner.

Such an IES is extremely difficult, perhaps impossible, to build, and is probably unnecessary to aim for. Current research is seeking ways that an IES may fall short of the 'ideal', yet lead to good learning.

One line of thinking reflects Papert's belief that exploration and building by the learner gives the best understanding: there should be no attempt to give the IES a sophisticated tutoring ability. It should answer the learner's questions, and perhaps make suggestions, but not set out to teach.

Contemporary cognitive science emphasises the diversity and richness of expertise. A human expert has a multi-level conceptualisation of the domain: there is factual information, the ability to use domain-specific methods and strategies, and to make judgements, and also the ability to monitor, control and evaluate his or her performance in the domain. A classical way to develop expertise is the apprenticeship method, in which the learner undertakes practical tasks, with demonstration, explanation, guidance, and criticism from an expert. Attention is directed not only to the learner's ability to do tasks, but also to his or her growing ability to plan, monitor and evaluate task performance.

Training or learning goals can often be regarded as the achievement of some level of expertise in the domain of interest. Viewing the target expertise in the rich, multi-level way described above, and recognising the value of the apprenticeship idea of 'do the task, with advice and evaluation' should give guidance to the design and use of IESs for training.

An important influence on prospects for applying IES research is the shortage of expertise in the field. It is, worldwide, a small research community, and there are only a few active research groups in the UK. If guidance and consultation is sought at the highest level, only a few experts would be potentially available to help. Some more specific examples of current research foci in IESs are described briefly below.

Intelligent help
Many software packages have an on-line reference manual for consultation by a user who needs help. An intelligent help system would monitor the user's actions, build up a picture of what the user did and did not know, and also what goal the user was trying to achieve. The help system could thus give more appropriate answers to a user's question, and could intervene to point out errors or suggest a more efficient way of achieving the goal. The user could thus work more effectively, and at the same time become more expert at using the computer application.

A full intelligent help system remains a big challenge, but even modest systems should be worthwhile. It is an attractive notion that software packages and computer-based machines should incorporate their own intelligent advice and training for users.

Direct manipulation
An extension to the learning environment approach is to provide objects on the screen that have properties similar to objects in the domain of interest. The user can manipulate the objects, build simulations, and 'learn by doing' in a computer-based representation of the domain. An example is ARK, the Alternative Reality Kit, which was developed in the *SMALLTALK* language by Randall Smith. In ARK an object created on the screen has many of the properties of physical objects, such as motion and momentum. ARK has been described as the best visual programming system yet; it will support simulation work in a variety of domains.

Simulate and coach
STEAMER and the RBT, the two large simulation systems described earlier, show that a simulation plus a flexible advising, helping and coaching function can be a valuable tool for learning. This approach can be extended. We could add a coaching function to a computer-based simulation, or to a machine itself - with an interface allowing the computer-based coach to communicate with the machine - or even to a direct manipulation or learning environment in which the learner is building and using his or her own small simulations.

Machine learning
The study of programs that learn from their own experience is an active research area. Any IES that adapts to the needs of a learner must learn about that learner, so machine learning is required. Machine learning ideas and techniques contribute widely in IES research, although they have usually been little noted other than by system designers. Machine learning is a difficult research area, but advances are likely to be valuable for IES development.

Domain independence
Many have dreamed of developing an IES shell, into which we would merely need to put a domain knowledge base to have an IES in the domain. Clancey linked GUIDON to three different knowledge bases and found that it could carry on discussions with a learner in each case. But the differing nature of the domain knowledge in the three cases meant that the discussions were of limited usefulness for learning. More recently Clancey has described the Training Express system, which is a less ambitious attempt to achieve some domain independence by giving the developer good tools to add some tutorial functions to any knowledge base expressed in a particular format. Training Express is now available commercially (see the annotated bibliography below.)

Achieving some degree of domain independence remains a goal of many IES researchers, but so far progress has been slow. Good tools are probably the most powerful domain independent things that AI work has given us to date.

4.4 SUMMARY

- Over some 30 years a wide variety of ways to use computers to support learning have been explored. Pioneering systems have demonstrated the application of fundamental AI ideas, including:
  - self-improvement of the system, by learning from interaction with learners;
  - separation of the knowledge base and reasoning functions;
  - building a learner model to allow adaptation to the needs of the individual;
  - limited natural language understanding;
  - use of an expert system as part of an IES;
  - giving of advice to the user of a computer or other device;
  - interactive simulation of a complex system; and
  - providing tools to allow the user to build and manipulate models.

- Cognitive science studies of human expertise emphasise the higher skills of judgement, problem-solving and control, as well as knowledge and skills.

- Wanting a computer system to be a full replacement for a skilled human tutor is generally a forlorn and unnecessary goal. But the list of features set out above shows that there are many ways that AI ideas can be exploited in computer systems intended to promote the development of some aspect of expertise.

5. AI APPLICATIONS IN TRAINING

Four approaches will be taken to discussing possible AI applications to training:
- how may the use of AI change the need for training?
- where in the traditional training process can AI contribute?
- what suggestions come from a conceptual analysis of training needs, and AI possibilities?
- what suggestions arise from IES research?

5.1 CAN AI TOOLS CHANGE TRAINING NEEDS?

Tools have been mentioned many times in the foregoing; use of AI tools is changing software engineering, and giving us more powerful and friendly applications, even on cheap microcomputers.

These advances naturally influence the need for training. For example, it may require some serious training effort for a novice to learn to make good use of the early word processing program Wordstar. By contrast, for a novice to learn MacWrite is scarcely a training issue at all, because the interface used is so much better. From a training point of view, perhaps the best applications of AI are those that reduce the need for having training at all!

It often happens that the detailed task analysis and knowledge elicitation that is needed at an early stage in the design of training suggests how redesign of the task, or provision of better job aids or interfaces may reduce the need for training, perhaps drastically. This has always been true, but becomes more likely now with the extra possibilities for task redesign and job aids offered by AI. In addition, more machines and processes now have a computer component, thus giving scope for intelligent help and advice to be incorporated.

Trainers should have, therefore, some familiarity with AI tools, techniques and possibilities.

5.2 AI AND THE SYSTEMS APPROACH TO TRAINING

The 9-step model of training described earlier may be regarded as the classic systems approach to training. Each step and sub-step in this process should be examined to see whether AI might be able to help, either by assisting the training designer, or by providing an alternative way to carry out a sub-task.

An expert system to assist

Where in the 9-step model might an expert system (ES) be used to help? The problem is one of identifying an area of expertise that is suitable for the building of such a system: it must not be too complex or amorphous - too hard for an ES - and not too trivial to warrant the ES development effort. It must also be an area in which human expertise is scarce or expensive, or again an ES would not be justified.

Examination of the steps of the systems model suggests that 'identifying the training need', and 'analyse the training need' are too broad and vague for ES treatment. Some aspects of 'design the training' and 'prepare the materials' have been proposed as candidates for ESs, but may be too limited or trivial e.g. selection of media, choice of screen layout, choice of question format. Some aspects of training design may, however, be appropriate: indeed the TC has already sponsored development of a Course Planner's Assistant, which reached early prototype stage. There has also been some work on Training Needs Analysis, although this seems a very difficult area for an ES. (See the appended list of references for pointers to further discussion of these possibilities.)

If a suitable area can be specified, it will be important to develop and use the ES as a support for the trainer, rather than to imagine that an ES will be able to replace the human training designer in any major respect.

If an ES can be specified, standard knowledge elicitation techniques, and ES development methodology will be needed. There remains debate, however, as to how best to proceed: should the software engineering practice of specify, then build, be used, or rather the AI practice of prototype, test and revise?

Awareness of AI ideas
The trainer following the classic systems model might benefit from familiarity with AI ideas and results. For example, recognition of the importance of knowledge elicitation and representation, and some understanding of AI techniques in these areas should help, even in the preparation of conventional training.

Enhancing CBT
It is an attractive possibility, and also in line with a gradualist, evolutionary approach, to consider how elements of IES functioning might be incorporated into CBT. Enhancement of CBT authoring tools is one promising approach. Even a modest ability to model and adapt to the individual user, or to interact in something approaching natural language, would represent useful progress.

An especially promising approach is to use an expert system not as a teaching device but as a resource, to be called on by the learner, perhaps partly under the guidance of CBT. An ES could provide the domain knowledge, to be explored with some CBT guidance, or an ES could advise a learner trying to make the most of CBT. A variety of ways of using ESs to enhance CBT deserve to be investigated.

It is worth noting, however, that adding AI will not magically turn a frog into a princess; nor will it make, for example, old video or film suddenly relevant and interesting.

5.3 AN ANALYSIS OF TRAINING DOMAINS AND AI TOOLS
The initial aim was that the seminar should try to identify promising applications of AI by stepping back and considering the broad nature of training domains, and corresponding AI formalisms. This turned out to be too ambitious a task, but some of the discussion can be viewed as addressing the issue.

Training domains
Considering training domains, a rough and subjective classification could be offered of the feasibility of building, within a 1 to 2 year time-scale, some sort of prototype AI-using training system for a selected topic area within the domain:

Bad bets Domains that look too difficult, e.g.:
- higher-level management skills
- manual dexterity (excepting computer interfaces)
- training needs analysis
- motivation
- aesthetics
- attitude change
- interpersonal skills.

Possibles Domains that look hard, but may be worth considering, e.g.:
- planning
- optimisation, given a metric
- heuristic problem solving
- principled design tasks
- advice embedded in software in a low-level language, e.g. C
- second language learning
- real-time processes having feed-back.

Good bets Domains within which it should prove possible to find a suitable topic, e.g.:
- algorithmic or well-specified problem-solving
- programming
- fault-finding
- financial management
- procedural real-time processes, e.g. monitoring
- CBT authoring
- simple mechanisms, e.g. robot arm
- advice embedded in software implemented in a high-level language.

AI formalisms
Considering formalisms used in AI for knowledge representation, the most promising might be considered to be:
- production rules
- objects
- frames
- prototypes
- nets.

Other formalisms, each with important areas of application, include:
- logics, both traditional and exotic
- conceptual graphs
- lattices
- repertory grids
- scripts.

AI tools
We should take for granted the basic system development environments now available, such as *LISP*, *PROLOG*, and object-oriented programming. and proprietary systems such as KEE and ART.

Beyond these, the most promising commercially available AI tools might be considered to be:
- browsers (allowing a user to explore a system)
- hyper-text (for example Hypercard on the Macintosh)
- simulation environments
- direct manipulation interfaces
- knowledge elicitation (including protocol analysis).

Other useful tools include:
- animation
- touchscreens
- interactive debugging
- reason maintenance systems
- object servers
- video servers
- expert system shells (but with reservations because of their limitations).

Conceptual mapping of training domains
One of the fundamental conclusions of AI research is that knowledge representation is at the heart of the AI problem, and that the formalism chosen to

represent knowledge must be appropriate for the domain, and for the purpose of the system being built.

The discussion of IES research given above confirms the importance of choice of knowledge formalism. It also emphasises that there are many kinds of teaching and learning activity that an IES may be designed to support, and, further, that the type of learning activity will influence what knowledge formalism is best. For example, an Anderson-style tutor for an algorithmic skill might be based on production rules, whereas a simulate-and-coach IES for the same domain might better use an object-based formalism.

In other words, we should consider not just the mapping from training domains to knowledge formalisms, but the more complex mapping from training domains and types of learning activity to knowledge formalisms. Changing the emphasis slightly, we could consider a mapping from training domains to types of IES, where specification of an IES requires choice of both a knowledge formalism, and one or more kinds of learning activity. However it is viewed, this mapping exercise would certainly be a valuable undertaking, but it remains as a challenge for the future.

5.4 LESSONS FROM IES RESEARCH

Rather than a conceptual analysis of the full range of training domains, a more pragmatic and immediate approach to identifying promising applications is to use existing IESs as models. From a training point of view, the most promising approaches include:

- Simulation, with or without coach. Simulation tools may help, although some, such as SIMKIT, which runs on top of KEE, are very expensive. One example would be a computer simulation of complex printing equipment: service technicians could work with the simulator to develop the subtle diagnostic kills needed to tune this equipment directly. When a new model of printer was introduced, the simulation could be adapted to cover the new model fairly easily.

  Another example is TRILLIUM, a simulation of copier interfaces developed by Xerox. This simulator is used to experiment with designs for new interfaces, and also for staff training. It is worth noting that one reason for this approach working well at Xerox is that the same computers are used for R&D and by operational departments.
- Modelling, or direct manipulation environments, possibly with advice or coaching. A general tool, such as ARK, could be used to implement a modelling tool tailored for the domain, to be used by the learner exploring the domain. Advice and guidance could be in traditional worksheet format, or come from a computer-based coach.
- Expert systems, used in a variety of ways. For example, a learner might use an expert system not for its rather rudimentary explanations, but to give a comparison with his or her own growing expertise; a critical attitude to the advice given by the expert system would be vital. Or the learner might use good AI-based tools to build or modify an expert system, thus being forced to express and organise his or her own understanding in the topic area. Or CBT or other traditional means could be used to guide a learner who explores an expert system in the domain of interest.

- Integration of advice and training functions, including intelligent help, into software and computer-based equipment. Training is thus to some extent integrated with use, and takes place at the workplace.
- Tutors for defined procedural skills that need to be learned and practised; Anderson's tutors provide a model.
- Natural language understanding by computer: although limited, this can usefully be added to assist interaction with e.g. a database, or a CBT system.
- Hypertext and browsing: there are many possibilities worth investigation. A promising one is the rapid prototyping of CBT.
- Enhancement of CBT design and preparation by incorporation of an expert system into the classic model for training.
- Use of AI tools and interfaces to assist the preparation and delivery of CBT.

6. IMPLICATIONS FOR THE TC PROGRAMME

Demonstrators
Good demonstrators are undoubtedly the key to a successful programme Industrial collaboration is needed, and careful selection and planning so that a successful demonstrator can be seen to make a clear and valuable enhancement to learning. The challenge is to find projects that are feasible within 1-2 years, that address a practical learning or training issue and that draw on IES research expertise.

Exploiting the lessons of IES research is a vital aspect: there are only a few research teams in the UK at the forefront of this research, and they, quite properly, put their primary effort into research projects with a longer time-scale than the TC can support. Some form of consultancy with leading research workers may be the best way to ensure that demonstrator projects are informed by the best possible advice.

Overseas experience
Several papers noted in the appended list of references describe AI-based systems developed in the United States specifically to enhance training. It would be worthwhile to investigate these and any other such systems.

Higher and further education
Most discussion has been in terms of training, but trends in training are to emphasise higher-level skills and expertise, thus making training more like education. Much of the discussion applies also to many aspects of higher and further education. The general strategy of using IESs developed for research as models for applications and demonstrators is a promising approach in education as well as in training.

Promoting awareness
Awareness and promotional activities should be based mainly on demonstrators or other applications that can be seen to perform. Direct comparisons of AI-enhanced training alongside conventional CBT or other training could be informative. There is also scope for more general educational activities, to raise preparedness in advance of future AI applications, and to help trainers benefit from general AI ideas, such as the importance of knowledge elicitation and representation.

Dangers of over-selling
Past over-selling and disillusionment about AI, and the occasional use of the AI label to add glamour to systems that embody no AI should lead to caution. It is important to find and support example applications that are likely to have a marked impact in some particular training setting.

TC support for AI applications
Experimental IESs have given concrete demonstrations of many ways that AI tools and techniques can contribute to learning. No existing or foreseeable system approaches the flexibility of a skilled human tutor, but systems have shown how individualised guidance and explanation can lead to enhanced learning. Practically useful improvements over what is possible with CBT and other traditional methods have been illustrated.

Application of the advances made by IES research has been slow in coming, partly because there is a shortage of specialist expertise. The large development effort needed by IESs has also been a problem, although improved tools now give scope for more rapid building of usable systems.

The aims of the Programme are worthwhile and practical, and amply justify TC support. Future support by the TC of applied research, demonstrators, dissemination and awareness activities about AI applications to learning is also justified, and, if carefully planned, would lead to worthwhile improvements in the efficiency and effectiveness of training and education.

6. AN ANNOTATED BIBLIOGRAPHY

Clancey, W. J. (1986) *From GUIDON to NEOMYCIN and HERACLES in twenty short lessons.* The AI magazine, August. 40–60;187. (A readable account of one of the longest-running and most important IES research projects.)

Clancey, W. J. and Joerger, K. (1988) *A practical authoring shell for apprenticeship learning.* Paper delivered at the ITS–88 Conference, Montreal, June. (Clancey describes the Training Express system, which is distributed by Teknowledge, Inc. Clancey's address is: Institute for Research on Learning, 3333 Coyote Hill Road, Palo Alto, CA, U.S.A. 94304)

Dear, B. L. (1987) *AI and the authoring process.* IEEE expert, Summer, 17–24. (A paper discussing how expert systems and AI tools can be used to enhance CBT.)

Kearsley, G. P. (Ed.) (1987) *Artificial intelligence and instruction.* Addison–Wesley, Reading, Mass. (A collection of recent papers describing IES systems, from the viewpoint of training; several come from the American 'instructional system design' tradition familiar to CBT designers.)

Merrill, M. D. (1987) *An expert system for instructional design.* IEEE expert, Summer, 25–37. (A paper describing 'ISD Expert', an expert system that is intended to incorporate knowledge of instructional system design and thus be able to assist the CBT author.)

Naughton, J. (1987) *Artificial intelligence: Applications to training.* 2nd ed. A research study for the Manpower Services Commission, by the Systems Group, The Open University. (An introductory and very readable report describing AI, training, CBT, and AI applications in training.)

O'Shea, T. and Self, J. (1983) *Learning and teaching with computers.* Harvester Press: Brighton, Sussex. (Mainly discusses learning at school level, but includes a historical account of computer use in education, and readable introductions to AI ideas and their application to support learning.)

Self, J. (Ed.) (1988) *Artificial intelligence and human learning.* Chapman and Hall: London. (A collection of recent research papers; technical, but illustrating the wide range of IES research.)

Wenger, E. (1987) *Artificial intelligence and tutoring systems.* Morgan Kaufmann: Los Altos, Calif. (Tells the story of 20 years of IES research. Somewhat technical, but comprehensive.)

APPENDIX A: ARTIFICIAL INTELLIGENCE APPLICATIONS TO LEARNING
– An update of the Training Commission's Programme

Since 1985 the Learning Technology Unit (formerly the Training Technology Section) of the Training Commission (TC) (formerly the Manpower Services Commission) has believed that the UK must come to terms with the application of AI techniques as an extension of existing and emerging learning technologies and methodologies. There has been a build up to a drive towards the "learning machine" – the intelligent tutoring system.

This drive led the TC into a number of projects very largely based on the application of Expert Systems to training and educational needs. Projects were funded to produce expert systems in:

- report writing;
- fault-finding and comparison with existing training methods;
- stock and inventory control for retail managers and supervisors;
- taxation for managers of small businesses.

A Further Education College was funded to produce a starter pack in the use of Expert Systems across the FE curriculum and to conduct training in the use of Expert Systems for trainers seconded from industry. Another ambitious project attempted to produce a generic training needs analysis system. Whilst a paper model of such a system was produced, the final outcome was a prototype course planner's assistant.

These projects gave an insight into what was possible and in August 1986 the TC agreed to fund a programme called "AI Applications to Learning". It would run for 3 years – April 1987 to March 1990 and have a total budget of £3.2 million; £0.6m. in the first year and £1.3m. in each of the next two financial years.

The overall aims of the Programme were and remain:

- to explore the use of AI techniques in developing more effective training methods;
- to accelerate the appropriate application of AI techniques in learning;
- to encourage UK industry to become more competent and competitive by providing evidence and demonstrations of more cost-effective means of training.

To help the TC derive a strategy, a wide cross section of experts were asked for advice and they gave an enthusiastic response. Based on this response an expert External Advisory Group was formed, this is some 25 strong and meets regularly to advise and to monitor the Programme. Based on such advice the first phase, which started in April 1987, had three elements:

- evaluation;
- demonstrator projects;
- survey reports.

Two parallel evaluations are running for the three years of the Programme and will provide for external evaluation of the impact of the Programme as a whole. They are both formative and summative and concentrate on the impact of the Programme on the Further and Higher Education sector and the Commercial and Industrial sectors respectively.

Survey projects were conducted and took the form of a mapping exercise of AI tools and intelligent courseware and a skills inventory, or training needs analysis, of the key people who could enable, influence or facilitate AI Applications to Learning. The surveys reported at the end of August 1987 to the TC and the External Advisory Group and helped frame the direction and criteria for the second phase of the Programme. An enhanced report will be available in September 1988.

Three demonstrator projects were selected from over £6 million's worth of proposals, 58 in all. They addressed learning needs in manufacturing industry, small businesses and the Further Education Sector. All proposals had to meet strict criteria and in the first phase three were selected:

- an expert system on customer complaints procedure for use by catering students and small businesses;
- computer based training and simulation linked with an expert system to train staff to operate and maintain process plant;
- an expert system training package to help owner-managers in business planning.

The External Advisory Group met in September 1987 to discuss the surveys and to advise on the way ahead. Following on from that meeting the TC decided the strategy for the second phase of the Programme.

Firstly, it was decided that more projects were needed to provide demonstrations of the application of intelligent solutions to specific learning needs. Consequently, Phase 1 was extended and a further four demonstrator projects were selected from the original 58 submissions. They were asked to resubmit to conform to additional criteria.

They covered a variety of areas of application:

- training in Statistical Quality Control;
- embedded operational support for naive users of a business software package;
- management training in Quality Management;
- courseware for decision makers in Information Structuring.

To increase the range of demonstrators still further, the TC called for further demonstrator proposals, conforming to tight criteria derived from the External Advisory Group meetings. From 78 proposals received, four were finally selected:

- an intelligent interactive video simulation on the fire control of large incidents for West Midlands Fire Brigade.
- intelligent courseware for insurance underwriting incorporating a trainee model. This will demonstrate the effectiveness of a unique AI-based authoring system designed specifically for training and compares the effectiveness of the AI approach to producing courseware with conventional CBT.

- an intelligent training and advice system for small businesses on employing people. This will also evaluate the Hypercard system in the production of intelligent computer-based courseware.
- two tutorial training programmes for the fluid power industry which will also evaluate an authoring system requiring no programming expertise and a system for guided discovery learning.

Project submissions were also called for which identify the key people who could enable, influence or facilitate AI applications to learning, map their skill needs and design and run pilot courses or courseware. In this way some of the skill shortages will be addressed whilst establishing a clearer picture of the needs of the original 25 submissions. One was selected which will develop a 5-day course to provide skills relevant to the development of Intelligent Learning System and Courseware for Managers and Trainers.

A sub-group of the External Advisory Group produced a specification for a major awareness campaign. This specification will form the basis for commissioning communication experts to design a campaign and the methods of implementation.

As a result of the recommendations made at a sub-group of the External Advisory Group endorsed by the main group, two further projects were funded:

- A two-day seminar to consider how AI can be applied to training by making a taxonomy of training domains and mapping knowledge representation so as to guide, inform and map the remainder of the Programme. [This report resulted from the seminar.]
- A survey to be conducted by a consortium which will inform and guide the rest of the Programme and will insure that proper account is taken of other initiatives and duplication of resources avoided. It will carry out a formal survey of research and development projects in the UK and overseas in which AI techniques will be used. Emphasis will be on AI Applications to Learning but will survey other developments which may eventually influence the education and training communities.

A two-page summary of each project mentioned in this brief update is available from:

Sharon Rubery
The Learning Technology Unit
Training Commission
Room W533
Moorfoot
Sheffield S1 4PQ

APPENDIX B: SEMINAR PARTICIPANTS

Dr M Barrett, Expertech Ltd

Mr R J Bland, Chartered Institute of Building

Dr G D A Brown, University College of North Wales

Dr G Cumming, University of Lancaster and La Trobe University

Dr J Gillingham, Training Commission

Mr J R Hartley, University of Leeds

Dr F Hickman, Knowledge Based Systems Centre Ltd

Professor R Lewis, ESRC InTER Programme, University of Lancaster

Mr J Lougher, British Steel Corporation

Mr T Mace, Artificial Intelligence Ltd

Ms A Mclure, Training Commission

Professor T O'Shea, The Open University

Mr R Rivers, Logica Cambridge Ltd

Mr J B Rynn, British Gas

Dr N Sharkey, University of Essex

Mr H Soltan, Cranfield Institute of Technology

Mr J Stock, Crosfield Electronics

Mr R Wormald, Training Commission

ESRC PUBLICATIONS RELATING TO INFORMATION TECHNOLOGY AND EDUCATION

Available from:

Economic and Social Research Council
160, Great Portland Street
London W1N 6BA – Telephone: 01-637 1499

Microcomputers in Education – a framework for research
(Morley Sage and David J. Smith)
Published by SSRC 1983 ISBN 0 86226 125 2 Price £2

Intelligent knowledge-based Systems
– UK Social Science Research Inputs
(Linda A. Murray and John T.E. Richardson)
Published by ESRC 1984 ISBN 0 86226 108 2 Price £2

Information Technology and Education
– signposts and research directions (Editor David J. Smith)
Published by ESRC 1984 ISBN 0 86226 152 X Price £5

From:

School Government Publishing Company
Darby House
Bletchingley Road
Merstham
Redhill, Surrey, RH1 3DN

Information and Communication Technologies:
Social Science Research and Training
Volume 1: An Overview of Research
William H Melody and Robin E Mansell
ISBN 0 866226 179 1

Volume 2: National Directory
Robin E Mansell (editor) assisted by Barbara Richards
ISBN 0 866226 184 8

Price £10 (two volumes)

From Blackwell Scientific Publications:

Trends in Computer Assisted Education – the proceedings of the 6th Conference on Computers in Higher Education held in April 1986 and supported by CET and the ITE Programme. Published January 1987.

Computer Assisted Learning in the Social Sciences and Humanities – the proceedings of the ESRC International Seminar held in April 1986. Published summer 1987.

[For ITE Programme and InTER Programme papers, see outside back cover.]

Papers currently available at no cost from:

ESRC – InTER Programme, Department of Psychology,
University of Lancaster, Lancaster LA1 4YF

Telephone:	0524 – 65201 Ext 4547 (24hr. answering)
Fax:	0524 – 63806
DIALCOM:	10001:yma027
JANET:	psg001@uk.ac.lancs.cent1
EARN:	psg001%uk.ac.lancs.cent1@ac.uk
DARPA:	psg001%lancs.cent1.ac.uk.@cunyvm.cuny.edu
EuroKom:	ESRCITE_Prog_Lancaster%eurokom.ucd.irl@euroies.uucp

ITE Series of Occasional Papers:

ITE/3/85	Spring Seminars Report, May 1985
ITE/14/86	International Electronic Services, November 1986
ITE/16/86	Information and Communication, December 1986
ITE/19/87	The ESRC – ITE Programme – two years on. May 1987
ITE/24/88	Teacher – fellow Research 1986/87 – abstracts, January 1988
ITE/25a/88	Research in Progress – research studentships, January 1988
ITE/25b/88	Research in Progress – research projects, January 1988
ITE/26/88	Learning through Microworlds – seminar report, January 1988
ITE/27/88	Authoring of Computer Based Training Materials – a study report to the Training Commission, June 1988
ITE/28/88	A series of four papers reporting on recent research of particular interest to: – primary teachers; – secondary teachers; – LEA Advisers; – senior policy – makers June 1988
ITE/29/88	Learners' concepts in mathematics and science – a seminar report, August 1988
ITE/30/88	Information Technology in Primary Schools – a research review for the Scottish Education Department, August 1988

InTER Series of Occasional Papers:

InTER/1/88	The Information Technology in Education Research Programme, August 1988
InTER/2/88	Artificial Intelligence Applications to Learning and Training – a seminar report to the Training Commission, August 1988

A full list of papers in the series is available from the Programme. Some not listed above may be available in photocopy form.