Journal of Courseware Engineering, Vol. 2. pp. 0-0, 1999-11-09

Printed in Norway

Seeing Learning as Work: Implications for Understanding and Improving Analysis and Design

Peter Goodyear

Centre for Studies in Advanced Learning Technology (CSALT)
Dept. of Educational Research, Lancaster University, LA1 4YL, England

Abstract – *It is becoming increasingly difficult to make clear separations between 'learning' and the rest of life. In this paper, I suggest that there are some good reasons for removing some of the 'otherness' about learning - for opening up new perspectives on analysis and design by casting learning as just another kind of work. There are some limits to, and flaws in, this proposition, but it does have the virtue of making us look at why user-centered approaches to technology development have not been very evident in educational technology. The paper includes some illustrations of what can be gained by 'seeing learning as work', drawn from the EU project SHARP.*

INTRODUCTION

This paper examines some important areas of convergence between the disciplines and practices of (a) educational technology and (b) performance improvement. Educational technology, for the purposes of this paper, is the field of theory and practice that focuses on the support of learning through the principled application of technology-based learning resources. I am particularly concerned with what are commonly called the 'upstream' activities in this field, notably the activities of analysis and design, and with learning resources that are computer based. These activities are acknowledged to be particularly important, since errors made 'upstream' flow down and affect all the stages that follow. Computer-based learning resources, in general, are particularly complex and problematic and require extra care and skill in their analysis and design (Pirolli & Greeno, 1988; Goodyear, 1997a). If the design is flawed, no amount of fancy programming will remedy things. I use the term 'performance improvement' to denote an emerging field of practice and theory which is concerned with methods of improving human performance in the workplace. Such methods may, but need not, include training interventions. They also include job redesign and workplace redesign.

My paper is concerned with both understanding the processes involved in the analysis and design of computer-based learning resources and *improving* those processes. For the last ten years I've been teaching postgraduate courses which are intended to improve the work of people who design computer-based learning resources and I've been engaged in a number of R&D projects which have tried to produce better tools for such people (e.g. Goodyear & Steeples, 1992, 1993; Goodyear, 1995, 1996). This concern for improvement introduces an element of 'indirection' into my paper, which complicates the presentation but also promotes the convergence to which I referred. That is, I'm interested in ways of improving the performance of people who work as educational technologists and among my methods are both (a) helping them learn to do things in better ways (b) helping retool their workplace. The educational technologists themselves are (of course) committed to improving the performance of other people (the learners whom they serve). Part of what I have to offer, through this paper, is a perspective on learning which casts learners as workers. I'm suggesting that educational technologists have much to gain by seeing certain forms of learning as 'just another kind of work'. Thus reducing the distinctiveness of what educational technologists are trying to do can open up possibilities of learning lessons which have been learned the hard way in areas like 'performance improvement'.

The Argument

My argument is a simple one. It is easiest to make in the context with which many of us in academia are most familiar - namely in relation to the learning environments we try to create for students in our university. Our central concern is to reach a better understanding of how to design and manage learning environments that are as supportive as possible of students who are working towards the educational objectives set for them. Information technology is playing an increasingly important part in such learning environments. We know from a mass of studies of success and failure in the implementation of IT systems, that they must be designed around a thorough understanding of the real needs and actual working practices of the intended users of the technology (e.g. Greenbaum & King, 1991; Schuler & Namioka, 1993). Systems that are designed around idealized models, and/or managers' versions of what the end-users' work actually consists of, usually prove to be expensive failures. It is reasonable to suppose that this is also true of IT systems which are designed for student users in higher education. Indeed, there is some evaluative evidence to suggest that educational software which is designed without a proper consideration of such issues will go unused by students (e.g. Draper et al, 1994). So the first part of the argument says that we should use best practice from the field of user-centered system design, participatory design etc. in developing the IT components of a higher education learning environment. We should take students' work seriously in designing such systems, and not make unfounded assumptions about the nature of that work.

The rest of this paper develops this argument and looks at some of its implications for understanding and improving the processes which are involved in the analysis and design of technology-assisted learning environments. In its latter parts it draws on some current work we are undertaking in the SHARP project, a goal of which is to harness asynchronous multimedia communications technologies for supporting the collaborative development of innovative working practices (see SHARP Project Team, 1998; Goodyear & Steeples, 1999).

THINKING OF THE LEARNER: IMPLICATIONS FOR ANALYSIS AND DESIGN

However implicit it may be, our 'model' of the learner influences our approach to analysis and design. One can think about learners in a number of ways. Here I'm going to use just one dimension, which I'll label 'highly compliant: highly autonomous'. Rather than make a general argument about whether we ought to regard learners as compliant or autonomous, I want to claim that (a) there are occasions when it is quite legitimate to expect learners to be compliant - to follow our instructions scrupulously, but (b) there are also circumstances in which it is unreasonable to want learners to be compliant - especially when we want them to take increasing control of their own learning, and (c) designing for compliant learners is (or should be) different from designing for autonomous learners. The same goes for analysis.

Analysis, Design and the Compliant Learner

Compliance is legitimate, sometimes. It is no coincidence that some of our central approaches to Instructional Systems Design (ISD) developed in the service of military training. Modern military problem-solving may well require some independence of mind. Nevertheless, the dominant conceptions in military training quite reasonably stress standardized procedures, well-defined target knowledge and compliance in both learning and learners. When an instructional designer is told very clearly by subject matter experts that there is just one good way of doing something, it makes their specification of aims, objectives and instructional events easier and more formulaic. A similar account can be offered in civil domains, especially in areas of high risk training for emergency situations, or training for repetitive processes. When training needs to provide a predictable response, ISD and a compliant learner are the order of the day.

Expecting compliant learners is reasonable in other educational settings: including some of the settings for which we design in higher education. For example, we still run laboratory classes in which student chemists are expected to learn safe procedures by working step-by-step through the instructions in a laboratory class handout. It's not just safety that may be at stake. There are times when we want students to practice accuracy in measurement or care in analytic technique. And organizing a learning activity for a group must sometimes rely on compliance: as individuals suppress their anarchic tendencies for the benefit of group outcomes. It is easy to think of additional examples, where it is quite legitimate and realistic to expect compliant learners, or at least tolerably compliant learners, and to draw on established methods of needs analysis and instructional design. We can put wholehearted effort into designing good tasks in the expectation that learners will do what we ask of them.

Consequences for Analysis and Design

The main consequence for analysis and design is that we can assume that learners will do what we tell them, or at least that they will try. A corollary is that they won't be surprised if the technology we produce restricts them to the space of our intentions.

Analysis, Design and the Autonomous Learner

There will always be obstreperous unpredictable learners. But sometimes - actually quite often - we *want* learners to show some independence. We value creativity, originality, self-reliance and abhor pedagogical spoon-feeding: or at least we say we do. I'm thinking particularly of higher education, but similar arguments apply elsewhere - whether in continuing professional development programs or reflecting-on-action in the workplace. Learning organizations ought not to expect compliant learners.

Valuing Autonomy

In higher education, there are a number of reasons to think that there is more than an empty rhetoric to our valuing of the autonomous learner. The massification of higher education, over the last 30 years or so, has shifted the nature of interest in autonomous learning, but there are few signs that it has reduced in importance. Students at Oxford may continue to regard lectures as optional and to fill their waking hours as they see fit. At less privileged universities, 'student-managed learning' may be a positive gloss on an economic necessity: resource-based learning in resource-starved settings. In any event, many universities are now constructing high-level strategies for teaching and learning based on an assumption that students will take increasing responsibility for managing their own learning (Vermunt & Rijswijk, 1988; Janssen, 1996).

Why Autonomous Learners Don't Do What We Tell Them to Do

There is a convincing body of evidence that students in higher education act (more or less) rationally in pursuit of a number of quite different goals. Becker et al. (1968) synthesized a very large volume of detailed ethnographic evidence to conclude that students have a rounded view of what maturation in the college environment involves. Academic goals are only part of the story. Furthermore, within the diverse set of academic goals offered to our students, cues are used intelligently to determine which goals are really important and which are merely espoused (Miller & Parlett, 1974). In short, for most students, most of the time, the assessment system plays a dominant role in shaping learning activity (Marton et al, 1997). Some students are very disorganized, others *are* strongly motivated by an interest in their academic studies, but the majority do a reasonable job of work: trying to find a realistic balance between the many calls on their time, focusing most of their efforts on the next piece of assessed work and finding little if any time for 'optional' or 'enhancement' learning activities (see e.g. Entwistle & Ramsden, 1983; Entwistle, et al., 1991; Draper et al, 1994). Autonomous learners are just like the rest of us - just plain folks, struggling to meet deadlines and relegating all that stuff which we know is good for us but which we'll just have to postpone for now. This has an important consequence for the development of computer-based learning software. No matter how breathtaking the multimedia or how pure the instructional design, computer-based learning resources will not be used (other than at the margins) unless (a) their use is mandatory, or (b) successful completion of an assessed task requires their use. It is important to recognise that this does not mean that students are Luddites. Far from it. Students seem to be making appropriate use of new technology when it can help them meet their goals. But they are mainly using the generic technologies that all just plain folks use: wordprocessing, email and browsing software, plus the specific 'tools of the trade' of their subject discipline or intended profession.

RETHINKING ANALYSIS AND DESIGN

How does that change the agenda for those of us who are concerned with designing technology to help students learn? Essentially, it simplifies the job. Our work falls into two neat parts. One part of it closely resembles the work of other people (outside education) who are interested in the design of technology that helps people at their work: the information systems designers, requirements engineers, human factors specialists and so on. In casting learning as work, the barriers come down between education and the rest of the world, and we can freely draw on, and contribute to, the knowledge bases and practices of all those other experts whose efforts we have habitually ignored. The second part of the work is special to us, however. It focuses on the design of good learning tasks.

Learning Tasks

I follow the French ergonomist Alain Wisner (1995, 597), in distinguishing between 'the prescribed work (the task) and the real work (the activity)'. This simple distinction, though a commonplace observation in everyday life, is sometimes overlooked – whether in the implementation of computer systems in offices or in the design of supportive technology for student learning.

A 'learning task', then, is what we ask learners to do. It is not necessarily the same as what they actually do (their learning activity). The analysis and design processes associated with creating good learning tasks for autonomous learners can be very similar to what we might think of as traditional analysis and design processes. That is, we still need the understanding of learning theory and the tools of ISD that we have always needed. But there are two significant differences. First, we are only in the business of designing tasks: not of embedding those tasks in some courseware. Second, we can usefully adopt an approach that might be called 'reflexive instructional design'. That is, our design of tasks can be informed by any knowledge we may have of the ways in which our autonomous learners transform task definitions into actual learning activity. I have not got space here for a fuller exploration of reflexive instructional design. The key point is that designing good learning tasks remains a core job for us.

User Centered Educational Technology

The big difference is in how we then approach the design of technology. Instead of insisting on a tight coupling between task and technology, which is the case when we hard-wire learning tasks into a courseware package, the aim should be to loosen the coupling and build supportive technology around actual activity rather than around prescribed tasks.

This is also a leap that is more radical at the level of philosophy than at the level of method. It *is* quite a shake-up to our ways of viewing the world: of our definitions of the roles of educational technologist and learner. But at the level of method, all it requires is that we go and find out what our learners actually do, and how we can improve the technology they use in doing it.

A nice example comes from the work of a former Lancaster colleague Mike Twidale. Mike Twidale's ARIADNE project (Twidale et al, 1995) had the goal of helping university students learn to search on-line databases, especially library databases. One approach to doing this, which we might (unfairly) call the traditional ISD approach, would set out to analyse the knowledge and skills needed for effective database searching, and design and develop some courseware which helped students acquire this knowledge and skill base. 'On-line searching techniques' is a dull topic so the courseware would need to be multimedia and motivating. It wouldn't run on the low tech library terminals, but would be just fine when installed on individually bookable multimedia PCs in a study center (where the students never go).

Mike's team had a better idea: they chose to start by seeing how students actually do their searching in the library. Within minutes they scored a major design insight: many students search collaboratively. The physical layout of the terminals is important to this. When a number of terminals are set up close together, one student can (and does) lean over to another and ask for help. A technology designed for individuals is actually used by twos and threes. Output designed for a 1 on 1 interaction (where much of the cognition is going on silently inside a single user's head) should actually be designed for a 1 on 2 or 1 on 3 interaction (where much of the cognition is distributed). Instead of designing some multimedia courseware, Mike's team designed a better interface. Instead of assuming this was a learning problem, they worked directly on performance.

This specific outcome is not a necessary consequence of 'seeing learning as work'. The general point is that it is important (and possible) to spend some time finding out what learners actually do. Understanding students' needs and practices is the key to user-centered educational technology.

Holistic Educational Technology

A related conclusion is that analysis and design need to engage with the totality of the learning environment, not just with isolated courseware components. An ergonomics of learning environments needs to attend to all of the significant relationships between the learner - as worker - and the physical environment which helps shape and sustain their activity. A number of commentators have raised our consciousness of the systemic nature of learning environments, for example by exploring some of the implications of ecological metaphors (see e.g. Hannafin & Hannafin, 1996). This presents serious challenges. It is intellectually difficult to construct a systemic understanding of (say) the nested set of learning environments within which specific communities of higher education students work. In addition, there is little history of efficient collaboration between the various agencies and actors responsible for designing and managing learning environments (see e.g. Ford et al., 1996). A first step on the path of better collaboration is a common language. I have a hunch that many who work in agencies on the 'hard' side of the learning environment (computing service personnel, for example) find the language of user-centered systems design less rebarbative than the language of pedagogy and ISD. A holistic, user-centered, educational technology may provide our best chance of bringing together the separate agencies and actors in a common operational framework. Seeing learning as work can help demystify what students do and what they need.

AN EXAMPLE: SHARP

In the last part of this paper I would like to give an example of a project in which we are trying to make practical progress with some of these ideas. The SHARP project is mainly concerned with programs of continuing professional development (CPD) at the postgraduate level. The approach is best described by looking at (a) 'high level pedagogy', (b) the kinds of tasks and activities that flow from this 'high level pedagogy' (c) the kinds of technological support suited to these activities.

High-level Pedagogy

We use the term 'high level pedagogy' to denote an important layer of design abstraction that sits between educational philosophy and action. (For example, one's philosophy may be constructivist; the high level pedagogy one selects or creates for a course is what gives expression to this philosophy; it is the source from which strategy and tactics can be derived.) Examples of high level pedagogy might be 'guided discovery learning', 'problem-based learning', 'programmed learning', or 'computer-supported collaborative learning'.

In SHARP, the high level pedagogy has no particular name (as yet) but brings together elements of cognitive apprenticeship, reflective practice and collaborative learning within a geographically distributed community of practice. It gives at least as much value to the tacit knowledge embedded in practice as it does to abstract, symbolic, articulated, research-based knowledge. It is particularly appropriate where CPD is taking place in an innovative, dynamic field: where much of the 'state of the art' is locked up in the practices of inventive individuals and their organizations. The high level pedagogy is intended to help unlock such knowledge, subject it to collective scrutiny and refinement, and feed it back into practice (SHARP Project Team, 1998; Goodyear & Steeples, 1999).

Tasks, Activities and Environment

The main tasks arising from this high level pedagogy are depicted in Figure 1, which can be thought of as a kind of learning cycle for members of a community of practice.

Insert correct figure 1 about here

SHARP has been focusing on the 'externalize' task. How can we get practitioners who are geographically remote from one another to create concise, vivid, mutually intelligible representations of their working practices, and of the knowledge embedded in those practices?

When we set such tasks, what activity results? And what kinds of technology are needed to support such activity? The short answer is that people respond to the task in a number of ways, engaging in different activities which result in different kinds of representation. The processes involved in creating these different kinds of representation vary, so that different configurations of technology are required. For example, in looking at the ways in which short videoclip representations of practice are created, we uncovered six different methods. (In what follows, the term 'subject' refers to the person(s) originating the material.)

1) **Fly on the wall**: In this type of representation, a camera is set up to capture practice in situ. The practice unfolds with as little attention to the video capture as possible. Depending upon the nature of the practice, the representation captured may or may not include speech from the subject(s) and also may or may not include visible physical activity. The practice may be primarily cognitive (i.e. going on inside the subject's head) and therefore may involve little visible action.

2) **Think Aloud**: This is also described as "talking while doing" or concurrent verbalization. Here the subject is filmed as they engage in the practice to be represented. The subject will have been prompted to think aloud as they work. The practice depicted in this type of clip is most likely to be one that is normally performed alone or at least to be one that does not normally involve speech (otherwise there is scope for interference and confusion between speech-in-the-task and speech-about-the-task). The difference between this type of representation and the following (action with commentary) depends on a claim about the validity of using concurrent verbalization to gain access to the subject's cognitive processes. Drawing on the views of Ericsson and Simon (1993) on this matter, we call a clip 'think aloud' if there are reasonable grounds to believe that what we are getting in the audio track is a reasonable reflection of the subject's task-related cognitive processes, and is not significantly biased by the subject's desire to provide a tidied up account of what they are doing.

3) **Action with commentary**: In this type of representation, the subject is encouraged to explain what they are doing as they work. This is different from thinking aloud because it requires the subject to create a commentary. They are creating their representation with an audience in mind, rather than being restricted to a concurrent verbalization on the task itself. The commentary may be a continuous explanation as the practice unfolds, or it may be that the subject begins, ends or interrupts the action with a commentary.

4) **Talking head**: This kind of representation moves further away again from the 'purity' of practice. Here the subject speaks to camera about a practice in which they engage and explains aspects of it. It is therefore likely to be a more discursive and reflective account than the accounts in the earlier examples. It may be done "on the fly" rather than rehearsed or it may be prepared in advance, though not scripted (see below). The subject will also be conscious of audience and will be explaining a practice rather than participating in it. The clip may offer a mixture of description and reflective comment. Of course as the subject talks about the practice they might wish to use artefacts to help clarify what they are saying. Such use of artefacts (etc.) from the work environment may be a useful source of 'grounding' in talking head clips: making it harder for the subject to depart (intentionally or unintentionally) from the actuality of their practices.

5) **Prepared script**: In this type of representation, the subject prepares a written description of (and possibly commentary on) a practice in advance of performance. This script may be used as a prompt, or followed closely for verbatim reporting. This is most likely to be a single subject activity.

6) **Acted**: In this final type, the subject engaged in a practice is played by an actor who adopts the role of the actual practitioner. This type of representation might involve scripting by the actual subject or could be prepared by a professional writer or production team in consultation with the practitioner. A familiar analogy would be with a training video produced commercially and used to illustrate how a skill or task should be undertaken, but the emphasis in the SHARP context would be on preparation of such materials in a lower key way for a small and known community of colleagues.

Asynchronous Multimedia Conferencing

The kinds of technology required for these different kinds of representation vary. An important part of the environment is the video capture technology itself - whether a professionally manned broadcast quality TV studio rig or a cheap monitor-top digicam or something in between. After that, the video clips need to be edited and probably annotated. The annotation may be through an audio track (such as an explanatory voice overlay), or through text annotations tied to temporary hotspots in the video image. Once this compound multimedia object, created as a representation of a selected working practice, is judged by the subject to be ready, it needs to be 'published' - placed in a shared space on the Internet, within which the learning community carries on its joint activity. Once published, the subsequent tasks on the learning cycle stimulate activities that need tools specialized for 'debate' and 'refinement' - mostly through various forms of annotation. A web of related multimedia objects is created and this becomes a shared resource for the community.

Most of the technology needed to support these activities exists, in a more or less primitive form. But there are few examples of environments that integrate them in a way that people who are not specialists in multimedia technology find sufficiently easy to use. We therefore see a need for an integrated technology that we call 'asynchronous multimedia conferencing' or AMC. AMC can be thought of as a multimedia enhanced version of the more familiar asynchronous text-based conferencing, provided by proprietary products like Lotus Notes or First Class. Such products do allow one to use multimedia files, but these have a kind of adjunct status. User-friendly AMC systems need to give a central place to multimedia objects and provide tools that are geared towards the editing and annotation of multimedia objects with other multimedia objects.

That was the short answer. If you are interested in longer answers please visit the SHARP website or read one of our papers (e.g. Goodyear & Steeples, 1999; Sgouropoulou et al, 1998).

CONCLUSIONS

'Seeing learning as work' creates an opportunity to demystify student learning. Not only can this have a beneficial effect on the multidisciplinary co-operations needed to design and manage complex, modern, technology-rich learning environments, but it also brings down some barriers separating educational technology from the practices of technology design in the rest of the world. This does not deskill us. We still need the best of what we know in order to design good learning tasks. And we can get closer to our learners: legitimizing their transformation of our tasks into their activities, finding out more about what they actually do, and what they need in order to do it better.

ACKNOWLEDGMENTS

I would like to acknowledge the support of the EC Socrates Open & Distance Learning Programme (grant number 40016-CP-2-98-1-GB-ODL-ODL) which has part-funded work on SHARP. My thanks also to members of the SHARP project team: Christine Steeples, Sonia Bartoluzzi, Nigel Oxley, Manolis Skordalakis, Tasos Koutoumanos, Cleo Sgouropoulou-Koutoumanou, Harald Haugen and Bodil Ask. For more about SHARP, visit the CSALT website at http://csalt.lancs.ac.uk/csalt

Finally my thanks to Mike Spector and his colleagues for inviting me to participate in the workshop on ‘Exploring the dimensions of performance improvement’, in Bergen, July 1998, and to Rita Richey and Dean Christensen for editorial care and patience.

REFERENCES

Becker, H, Geer, B & Hughes, E (1968) *Making the grade: the academic side of college life*, Wiley, New York

Draper, S et al (1994) *Observing and measuring the performance of educational technology*, TILT project report No. 1, Robert Clark Center, University of Glasgow, Scotland.

Entwistle, N, Meyer, J & Tait, H (1991) Student failure: disintegrated patterns of study strategies and perceptions of the learning environment, *Higher Education* (21) 249-261.

Entwistle, N & Ramsden, P (1983) *Understanding student learning*, London: Croom Helm

Ford, P., Goodyear, P, Heseltine, R, Lewis, R, Darby, J., Graves, J, Sartorius, P, Harwood, D & King, T (1996), *Managing change in higher education: a learning environment architecture,* The Society for Research into Higher Education/Open University Press, Buckingham

Goodyear, P. (1995) Situated action and distributed knowledge: a JITOL perspective on electronic performance support systems, *Educational and Training Technology International*, (32,1) 45-55

Goodyear, P. (1996) Asynchronous peer interaction in distance education: the evolution of goals, practices and technology, *Training Research Journal* (1) 71-102

Goodyear, P (1997a) Instructional design environments: methods and tools for the design of complex instructional systems, in Dijkstra, S, Seel, N., Schott, F., Tennyson, R (eds*) Instructional design: international perspectives,* pp83-111, New Jersey, Lawrence Erlbaum Associates

Goodyear, P. (1997b) The ergonomics of learning environments: learner-managed learning and new technology, in *Creacion de materiales para la innovacion educativa con nuevas tecnologias*, pp7-17, Malaga, Instituto de Ciencias de la Educacion, Universidad de Malaga (available in English and Spanish at http://www.ice.uma.es/memorias/congresos/Edutec97/edu97_co/edu97co.htm)

Goodyear, P & Steeples, C. (1992) IT-based open learning: tasks and tools, *Journal of Computer Assisted Learning* (8), 163-176

Goodyear, P. & Steeples, C. (1993) Computer-mediated communication in the professional development of workers in the advanced learning technologies industry, in Eccleston, J., Barta, B., Hambusch, R (eds) *The computer-mediated education of information technology professionals and advanced end-users*, Elsevier: Amsterdam, 239-247

Goodyear, P & Steeples, C (1999) Shareable representations of practice: asynchronous multimedia conferencing in continuing professional development, *Distance Education* (20, 1) 31-48

Hannafin, K & Hannafin, M (1996) The ecology of distance learning environments, *Training Research Journal* (1) 49-70

Janssen, P. (1996) Studaxology: the expertise students need to be effective in higher education, *Higher Education* (31) 117-141

Sgouropoulou, C, Koutoumanos, Goodyear, P, & Skordalakis, E (1998) A web based system for asynchronous multimedia conferencing, *Proceedings of WebNet'98*, Orlando, Florida.

Marton, F, Hounsell, D & Entwistle, N, eds, (1997) *The experience of learning: implications for teaching and studying in higher education*, Scottish Academic Press, Edinburgh.

Miller, C & Parlett, M (1974) *Up to the mark: a study of the examination game*, Society for Research in Higher Education, Guildford.

Pirolli, P & Greeno, J (1988) The problem space of instructional design, in Psotka, J, Massey, L, & Mutter, S, eds, *Intelligent tutoring systems: lessons learned*, Hillsdale New Jersey: Lawrence Erlbaum Associates.

SHARP Project team (1998) *Shareable Representations of Practice: An introduction to the SHARP project*, Lancaster, (available at http://www.lancs.ac.uk/users/edres/research/sharp/publications.htm)

Twidale, M, Nichols, D, Mariani, J, Rodden, T & Sawyer, P (1995) Supporting the active learning of collaborative database browsing techniques, *Association for Learning Technology Journal* (3,1) 75-9

Vermunt, J & Rijswijk, F (1988) Analysis and development of students' skill in self-regulated learning, *Higher Education* (17) 647-682.

Wisner, A (1995) Understanding problem building: ergonomic work analysis, *Ergonomics* (38,3) 595-605.

Figure 1