Developing task design guides through cognitive studies of expertise

In Bagnara, S. (Ed.) (1999) Proc. European Conference on Cognitive Science(ECCS99) Sienna, Italy (pp. 401-410)

Thomas C. Ormerod
Department of Psychology
Lancaster University
LA1 4YF, UK.
T.Ormerod@lancaster.ac.uk

James Ridgway
Department of Education
Durham University, UK
Jim.Ridgway@durham.ac.uk

ABSTRACT

This paper describes two empirical studies that informed the development of a Task Wizard, a computer-based system for presenting task design guidance. Two empirical studies compared specialist task designers with experienced teachers using task design observation and task sorting paradigms. In Study 1, designers worked initially in depth and then worked in breadth before returning to depth-first development, unlike teachers who moved from working in breadth initially to working in depth. In Study 2, designers' classifications of tasks were less dominated by deep conceptual topic-based categories than were those of teachers. These results are the reverse of expert/novice differences identified in the literature and have important implications for theories of design expertise in particular, and complex problem-solving and expertise in general. The studies also provide data on task classifications, evaluative criteria, examples of good and bad designs, and design strategies and practices that informed the development of task design guides. The guides are delivered in the form of a Task Wizard, a computer-based platform for supporting production, adaptation and selection of tasks for classroom activities and examination questions.

Keywords

Expert/novice differences, instructional task design, conceptual representations, structured and opportunistic problem decomposition, protocol analysis, card sorts, computer-based design guides.

INTRODUCTION

Many national education systems are attempting reforms that set out a range of ‘new’ educational goals, where students are expected to apply academic knowledge to real-world contexts, communicate effectively and work collaboratively. Thus, there is a large demand for new styles of task in education, in turn creating a need for effective task design procedures. We set out to investigate the processes underlying the development of tasks for instruction and assessment as part of a larger project exploring the nature, acquisition and support of task design expertise, funded by the UK Economic and Social Research Council. By a task, we refer to a relatively short and self-contained student activity that introduces, rehearses or assesses a specific curriculum objective or topic (an example is shown in Figure 1). Thus, our focus is at a finer grain size than researchers who investigating curriculum and training design (e.g., Perez, Fleming Johnson & Emery, 1995).

Task design can pose great difficulties for teachers, and it is apparent in both Mathematics and English as a Foreign Language (EFL) teaching (our chosen domains of study) that domain knowledge is not sufficient to guarantee skill in task design. It is also clear, where guidance in the process of task design is provided, that such guidance has no empirical basis. At the same time, task design presents an ideal domain in which to carry out research into problem-solving in general, and design expertise in particular. Task design is creative and yet constrained, and the scale of design problem makes it relatively easy to study. Specialist and non-specialist designers are available for study, and there is an enormous population of end users (e.g. teachers textbook authors, exam boards). The development of curricula and computer-based courseware has been studied extensively, but we are aware of no research that investigates specifically the process of task design. This paper reports a project to investigate the process of designing tasks for instruction and assessment in a systematic way, in order to develop task design guides that support the task design process. We focus in particular upon two empirical studies that contribute to the development of a computer-based task design guide disseminated across the World Wide Web (WWW).

Two findings are ubiquitous in the study of expertise. First, experts are able to call upon deep conceptual representations of domain knowledge, which they adapt to meet current task demands. In contrast, novices rely upon shallow representations that focus upon superficial features of the domain. Where expert

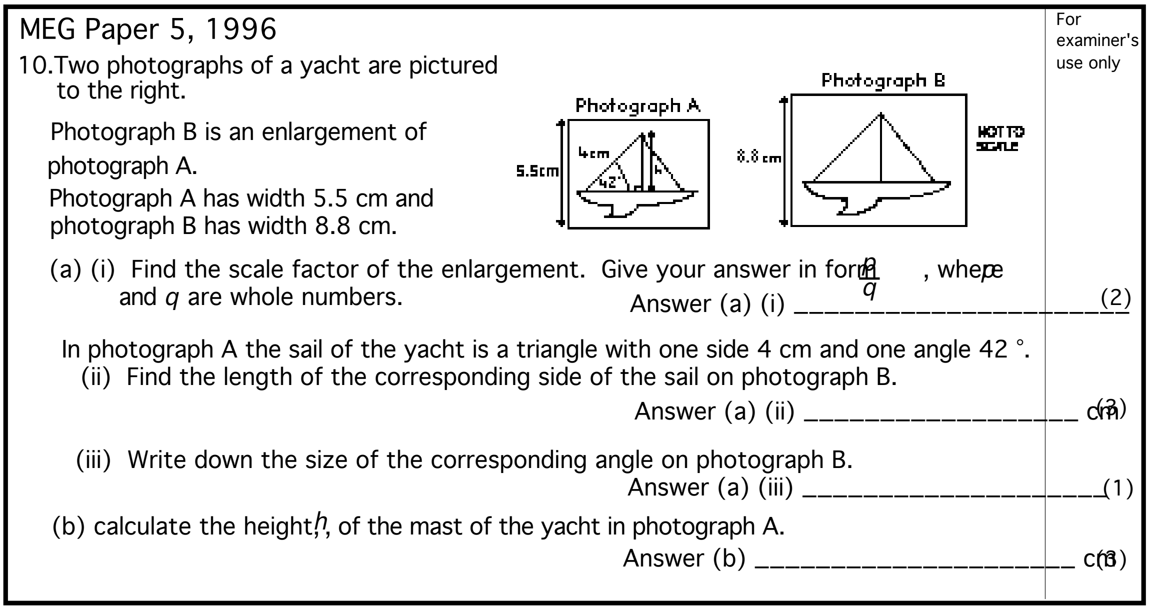

MEG Paper 5, 1996

For examiner's use only

10. Two photographs of a yacht are pictured to the right.

Photograph B is an enlargement of photograph A.
Photograph A has width 5.5 cm and photograph B has width 8.8 cm.

(a) (i) Find the scale factor of the enlargement. Give your answer in form $\frac{p}{q}$, where p and q are whole numbers.

Answer (a) (i) ______________________ (2)

In photograph A the sail of the yacht is a triangle with one side 4 cm and one angle 42 °.

(ii) Find the length of the corresponding side of the sail on photograph B.

Answer (a) (ii) ____________________ cm (3)

(iii) Write down the size of the corresponding angle on photograph B.

Answer (a) (iii) ______________________ (1)

(b) calculate the height, h, of the mast of the yacht in photograph A.

Answer (b) ______________________ cm (3)

Figure 1. An example of a GCSE Mathematics task used in the study (©MEG examination board, UK).

performance is based on deep representations, it is characterised by the rapid recognition of problem states and the structured development of solutions following a predetermined pattern (e.g., Chi, Feltovitch & Glaser, 1981; Koedinger & Anderson, 1990).,

Second, experts manage complex problem spaces, using the structured breadth-first decomposition of complex problems into manageable chunk sizes (e.g., Jeffries, Turner, Polson & Atwood, 1981). This observation is consistent with the view that breadth-first decomposition is more prescriptively optimal than depth-first decomposition (e.g., Wirth, 1971). A breadth-first approach maximises the exploration of the problem and minimises early commitment to solutions that might turn out to be unworkable when new aspects of the problem are explored. Breadth-first contrasts with depth-first decomposition, in which a specific sub-goal is pursued in detail to completion before the rest of the problem is explored. Depth-first decomposition has been found in novices' problem-solving performance (e.g., Jeffries et al, 1981), reflecting the fact that it generates a lower cognitive load than a breadth-first strategy since there is no requirement to remember explored but incomplete problem sub-goals. Experts may be able to alleviate the problems of cognitive load in breadth-first design by retrieving conceptual knowledge that structures the recall of problem sub-goals.

Deep conceptual representations and structured decomposition strategies confer many advantages, but there are situations where they can actively impair performance. For example, Wiley (1998) reports a study in which priming of domain knowledge impairs expert performance. These studies encourage a view that experts are unable to 'turn off' their deep domain knowledge when it is inappropriate.

Of particular interest to the current research is Wiley's suggestion that domain knowledge can inhibit creative problem-solving. Design is a creative problem-solving activity, and a large body of evidence has accumulated showing the same kinds of conceptual representation underlying expert design found in other expertise. However, when originality of output is the primary concern, prior knowledge may be less efficacious, perhaps even leading to 'design fixation' (Jansson & Smith, 1991). Similarly, the role of structured decomposition strategies in creative domains has been questioned. Guindon (1990) suggests that expert design is best characterised as opportunistic rather than structured, and that opportunism is an important source for creative ideas. The studies reported below, as well as providing data for developing task design guides, allow us to assess the role of conceptual representations and decomposition strategies in the creative design domain of task design.

STUDY 1: TASK DESIGN OBSERVATION

We recently carried out a study in which we analysed the verbal protocols of participants developing novel tasks to appear in English as a Foreign Language (EFL) textbooks. The study addressed two main questions. First, do task designers use problem-solving strategies similar to those of experts in other domains? Second, do specialist task designers and teachers differ in terms of their focus of activity during a design session? For example, Chi, Feltovitch & Glaser (1981) report that a major in physics problem solving was that experts tended to undertake more evaluation than novices. Similarly, Schoenfeld & Herrmann (1982) found an increase in monitoring and scheduling activities among expert Mathematics problem-solvers.

Method

Eight designers, experienced authors of task-based EFL texts, and 13 experienced EFL teachers with little or no task design experience undertook a two-hour session to produce a task that involved specifying a communicative activity for practising descriptions of people. Verbalisations and written notes were recorded and time-stamped. After transcription and segmentation, the resultant protocols were coded using a scheme defined *a priori*, in which the cognitive activity underlying each segment (e.g., evaluate, generate, monitor, schedule) and the referent (e.g., brief, task ideas, detailed design) was coded. Coding was carried out by two researchers who coded the transcripts independently. Agreement in cognitive act assignment was 76%, and in referent code assignment was 88%. While these rates may appear relatively low, this is to be expected for such an unconstrained problem-solving domain. Differences tended to be in acts involving cognitive control (e.g., differences in the use of Monitor and Reflect categories), and in grain size of the referent inferred from the protocol texts. Subsequently, the researchers met to discuss differences, and all bar 2% were resolved.

The coded protocols were then examined for phases of activity, distribution of activities, and adherence to decomposition strategies. Scripts were written in FileMaker Pro™ to count the frequencies of cognitive acts and referent types in each protocol, and the distribution of these across phases of activity. Phases were defined by time quartile and also by qualitative analysis of each transcript to determine major changes in designer activity. The extent to which designers adhered to structured decomposition strategies or demonstrated evidence of opportunistic decomposition was also assessed. This was done by constructing hierarchies of design referents for each protocol, and counting the transitions between referents that conformed to breadth-first and/or depth-first traversal.

Results and Discussion

Inspection of the designers' protocols for boundaries between phases of design activity revealed four distinct and ordered phases; idea generation, design description, expansion, and implementation. These are analogous to the stages of requirements specification, conceptual design, detailed design and implementation found in engineering, industrial and software design (e.g., Eason, 1988). The phase of idea generation, while it included activities devoted to understanding the brief, also was when designers searched for, rapidly evaluated, and subsequently accepted or rejected task ideas. Counts of referent transitions within each phase show that designers worked mainly depth-first during idea generation, 28.7% of transitions violating breadth-first (BF), 3.4% violating depth-first (DF), with 0% violating both structured approaches (see Table 1 for data for other phases). Breadth-first development increased during description and expansion, and depth-first development dominated final implementation.

The designers' approach can be characterised as follows: In the phase of idea generation, they adopt a generate and test strategy in the search for creative task

Table 1. Percentage of referent transitions that violate structured decomposition strategies (BF = breadth-first violations, DF = depth-first violations, BF&DF = violations of both) in each phase of design activity (*only 2 teachers showed an expansion phase). Note that the lower the percentage of violations, the more consistent is design with each decomposition strategy. Characterisation describes the overall design approach adopted in each phase by each group.

DESIGNERS

	Generate	Describe	Expand	Implement
BF	28.7	13.2	10.9	50.5
DF	3.4	16.1	18.7	1.9
BF & DF	0.0	5.2	1.6	4.9
Characterization	Generate & test	Balanced design	Balanced design	Completion

TEACHERS

	Generate	Describe	Expand*	Implement
BF	1.3	11.8	64.5*	25.2
DF	43.7	23.4	6.1*	5.3
BF & DF	0.0	2.8	0.0*	1.7
Characterization	Plan retrieval	Exemplar adapt	(Completion)	Completion

ideas, working in depth to rapidly explore and evaluate ideas ("fail fast and fail often"). In the phases of description and expansion of a single task idea, they switch to a breadth-first mode, the optimum strategy for ensuring that all design components receive sufficient development and evaluation without prematurely committing the designer through overly-detailed design work. This has been described as 'balanced design' (Adelson & Soloway, 1985). Working in breadth entails the scheduling of activities to be completed later. In the stage of implementation, all the contingencies among task components have been explored and evaluated, and so it is pragmatic again to work in depth.

The teachers showed only a short phase of idea generation, in which a single generic task was retrieved and then described, with little or no expansion prior to implementation. Counts of activity transitions show that they worked mainly breadth-first during idea generation. Depth-first development increased during description and expansion, and dominated implementation. The protocols of teachers suggest that task development was preceded by the retrieval of a task plan. In many instances participants stated a plan explicitly as a definition of what constitutes a good task (e.g., “warm up -> language input -> group activity -> denoument -> feedback”), and then retrieved and subsequently modified a task exemplifying this plan. Specification of the plan and retrieval and description of the exemplar gives rise to the kinds of breadth-first development seen in other plan-based design domains, such as routine programming (e.g., Soloway & Erhlich, 1984). Once retrieved, task implementation was completed in depth.

Analysis of referent transitions allowed us to test a hypothesis that skilled designers in a creative domain would show less adherence to systematic design practices than designers in highly constrained domains such as industrial and engineering design. Referent transition counts show that this is simply not the case. On average, each transcript contained approximately 400 transitions, only one participant making more than four structure-divergent transitions – there was little general evidence of opportunism. This is not to say that these transitions were unimportant, simply that they did not dominate design activity, a finding that challenges many contemporary design theories (e.g., Guindon, 1990; Visser, 1994). It is particularly interesting that

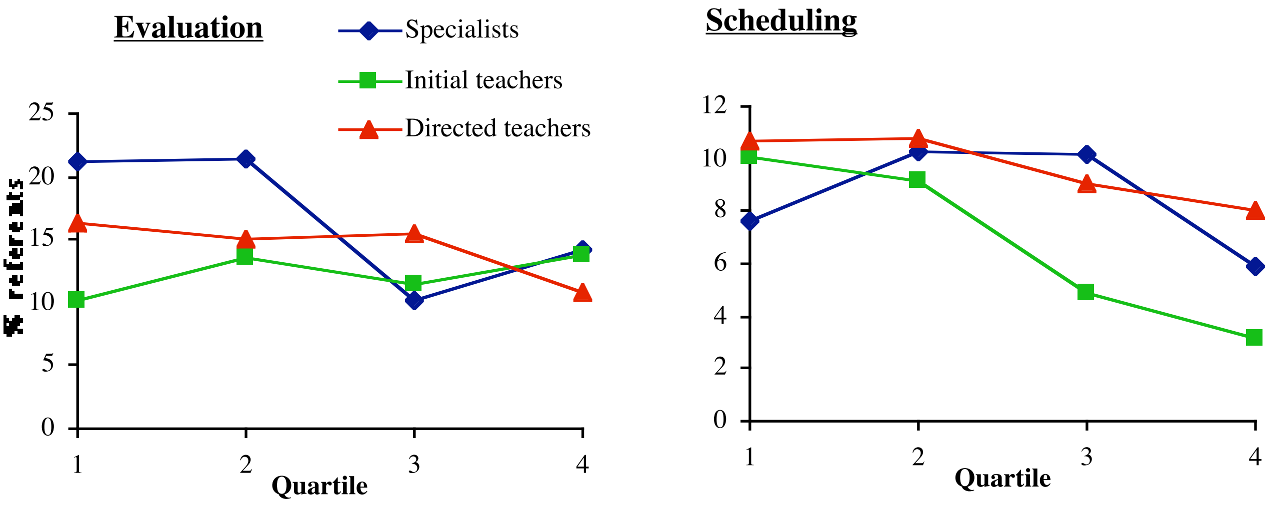

Figure 2. Mean % of cognitive acts involving evaluation and scheduling in Study 1

structure convergence is not a function of specialism, since researchers in other domains (e.g., software design) have found that experts are much more structure-convergent than novices.

Designers and teachers also differed systematically in the distribution of cognitive acts across time. Figure 2 shows the percentages of each class of cognitive act for each of four equal time quartiles. Statistical analyses of these data show no overall difference in the amount of evaluation or monitoring. However, significant interactions between time quartile and group were found with respect to evaluation, $F(3, 36) = 2.67$, $p<.05$, and scheduling, $F(3, 36) = 3.99$, $p<.01$. It appears that designers carried out more early evaluation than teachers and scheduled more than teachers during the middle phases. This can be explained by reference to the amount of activity congruent with breadth-first design which increased in description and expansion phases for designers. It declined across all stages for teachers, parallel with a decrease in scheduling across phases. This is consistent with setting up of breadth-first requirements through plan retrieval and their gradual resolution with increasing depth-first development. Thus, scheduling appears related to maintaining breadth-first decomposition, while evaluation seems to relate, at least in part, to the depth-first elaboration of specific details or testing of putative task ideas.

Returning to the first question addressed in the study, it appears that specialist task designers use decomposition strategies found in other domains. Approximately 98% of transitions conformed to a top-down structured strategy – there was little evidence of opportunism. Yet, in contrast to the teachers, the designers produced some highly original and innovative solutions. The large differences between the groups may, at least in part, be attributed to the fact that each group operated under a different key constraint. In the case of the teachers, this appears to have been the production of a usable task. It may be argued that retrieval of task types from memory is the best strategy to adopt under this constraint, since recalled tasks are in some sense warranted by the experience of the teacher or of others. Adding a demand for originality makes reliance upon retrieval less appropriate. The designers' "fail fast-fail often" strategy provides a mechanism for satisfying the originality constraint. Once an idea has been determined, they then manage the production of a practicable task by switching to a balanced design strategy. The key point is that, depending upon whether the participant is designing from scratch or on the basis of a retrieved conceptual representation of what constitutes a good task, the optimality of each approach is determined by the designing context.

STUDY 2: TASK SORTING

The results of Study 1 suggest that teachers' designs were constrained by application of a pre-determined task structure. This is consistent with Wiley's (1998) contention that conceptual representations can constrain creative problem-solving. So, what is the source of creativity among the specialist designers? It may be that designers do not have these conceptual representations, though we find this unlikely given that our groups were matched for teaching training and experience. Alternatively, designers may acquire strategies to bypass conceptual knowledge or acquire alternative conceptual structures that take precedence. In order to explore this issue further, we carried out a study using a card sort method (a more detailed report of this study is given by Ormerod, Fritz & Ridgway, 1999). The aim was to see whether designers and teachers differed in the types of dimension, assignment to categories under each dimension, and the order in which sort dimensions were produced.

Table 2. Percentage of designers (D) and teachers (T) producing sort dimensions , and mean position in which the dimension arose in the sort order. Other dimensions produced by more than one participant (% < 25 in each case) were Ramping, Complexity, Exam board, Wordiness, Response type, and Mark.

Participant	Math topic	Level	Openness	Thinking vs rote	Difficulty	Context	Prefs / turn-offs	Graphics
% D	100	60	50	14	14	33	29	14
% T	95	20	10	37	35	15	20	25
Position D	1.6	1.6	2.7	2.0	2.5	3.2	3.5	3.0
Position T	1.0	2.5	3.5	3.3	2.8	2.7	2.8	2.4

Method

Twenty Mathematics teachers and 14 specialist designers carried out repeated sorts of a set of 12 tasks (an example is shown in Figure 1) taken from a GCSE-level Mathematics examination paper (equivalent to the middle of high school), sorting the tasks into categories under dimensions that were meaningful to them. Participants were encouraged to sort repeatedly, so that the order in which sort dimensions were generated could be examined as a measure of conceptual precedence. The resulting sorts were then subjected to multidimensional scaling and cluster analyses.

Results and Discussion

A number of interesting differences emerged. Table 2 shows some of the common dimensions produced during the sorts. Designers produced reliably more sort dimensions than teachers, $t=2.17$, $p=.03$. Also, teachers produced Maths topic sorts earlier than designers, $U=85.5$, $p=.048$. Teachers produced more specific Maths topic dimensions than designers ('linear inequalities', 'fractions', 'number patterns' vs. 'Algebra and number'). Teachers & designers sorted similarly under most GCSE Attainment Target (Applying & using math; Algebra & number; Shape & space), though designers used the Data handling category more than teachers.

The results suggest that, while both teachers and designers have knowledge representations that allow them to identify core Mathematics principles underlying the tasks, the designers place less primacy upon them. Designers use a wider range of conceptual representations, and appear to suppress topic-based dimensions in favour of others (e.g., 'openness'). In general, the specialist sorts tended to be more superficial than the non-specialists. This result is in precisely the opposite direction other expert/novice sort studies (e.g., Chi, Feltovitch & Glaser, 1981), in which conceptual depth is generally associated with expert groups. Again, it is possible to account for these differences in terms of the optimal approach to undertaking different types of realistic task activity. Conceptual representations of Maths topics are precisely what is needed for selecting appropriate exercises to meet curriculum goals, which explains their primacy in the sorts of teachers. Task design, on the other hand, requires different kinds of knowledge to task teaching. 'Superficial' sorting dimensions may reflect the very things that make tasks interesting, original and practicable, which explains their primacy in designers' sorts.

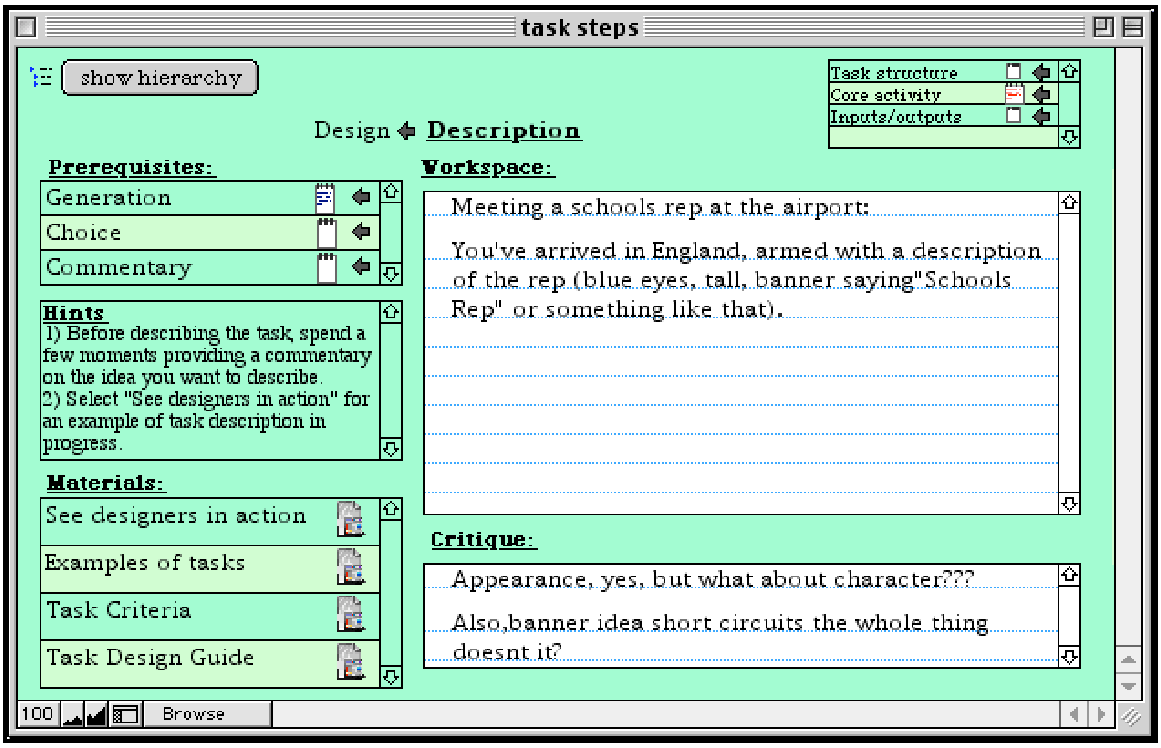

1) Before describing the task, spend a few moments providing a commentary on the idea you want to describe.
2) Select "See designers in action" for an example of task description in progress.

Meeting a schools rep at the airport:

You've arrived in England, armed with a description of the rep (blue eyes, tall, banner saying"Schools Rep" or something like that).

Appearance, yes, but what about character???

Also,banner idea short circuits the whole thing doesnt it?

Figure 3. A Task Wizard module. "Show hierarchy" accesses the module hierarchy, from which the user can navigate freely. The top-right box indicates sub-modules to which the user can navigate when required. The top-centre title is a clickable button for moving up the module hierarchy. The left panel shows prerequisites, hints and materials. Icons provide visual indicators as to whether these modules have been pursued. The icons next to items in the materials are URLs to web pages.

THE TASK WIZARD

The applied aim of our research was the synthesis of task design guides, which entailed three main activities. First, a literature survey was carried out in each domain, in order to classify task types and to assess existing task classification and design metrics. In the EFL domain, this was largely achieved through surveying published textbooks. In the Mathematics domain, our work was strongly influenced by the work of the Mathematics Assessment Research Service (MARS) funded by the National Science Foundation in the USA. Second, semi-structured interviews were conducted with experienced designers in each domain. These were analysed discursively, and commonalities and differences between the task design and classification approaches described by participants in each domain were noted. Third, and most importantly, the data from the empirical studies described above were used to inform the content and structure of the guides.

The guides provide prescriptive advice on task design, but are mainly geared towards providing examples of design practice observed in our studies, both good and bad, from specialists and non-specialist designers. We initially developed paper-based versions of the guides. However, there are a number of advantages to developing computer-based versions. For example, it allows the embodiment of a flexible procedural model of task development. This is important in the light of the evidence from Study 1 showing sophisticated changes in problem decomposition strategy over design phase.

A computer-based Task Wizard was created using Filemaker Pro™ (this platform was chosen in part because much of our data were coded in this form). The concept borrows from existing text and graphic generation aids (e.g., Microsoft's Chart Wizard). The Task Wizard supports the process of creating a task by guiding the user through of a set of hierarchical work modules, providing module-specific advice where appropriate. Each module provides work and critique spaces (the latter reflecting our observation that on-line evaluation, especially early evaluation of task ideas, is an essential aspect of task development). The user moves through the hierarchy as they wish, thereby supporting a mixture of design decomposition approaches.

Figure 3 shows an example screen from the Task Wizard. A list of task development materials is provided. This consists of WWW links to video clips of design specialists in action, critique and task generation transcripts in FileMaker Pro™, relevant sections of the paper-based guides, a list of task criteria that contains links to transcript examples, and sets of task exemplars. The use of links to web pages containing guide materials means that design guide materials can be updated remotely, neatly avoiding problems of legacy software. At the end of a session, the user can obtain an organised record of their task development activity, in the form of a computer text file. In principle, this could be extended to producing ready-formatted materials (e.g., teachers' notes, student worksheets, examination papers).

The Task Wizard embodies a flexible design cycle based upon our empirical observations of design phases, and also supports switches between structured decomposition approaches by allowing the user to traverse the module hierarchy in a self-determined pattern. The system can be configured such that each module has pre-requisites, that is, work modules that in an ordered design process come before the current module, as observed with the specialists in Study 1. A 'Hints' box provides module-specific advice to the user, derived from assessing the current module and any specified pre-requisites. For example, in Figure 3, the user is beginning to work on describing a task idea that was selected in the generation module. The Task Wizard recognises that the workspace for the generation module contains notes, but that 'choice' and 'commentary' modules have not been pursued. A hint to switch to the 'commentary' module is provided, along with a hint to look at a video clip of a design specialist carrying out a task description activity (sampled from the data of Studies 1 and 2). The Task Wizard provides a seamless environment for the following activities:

- Developing a task from scratch: This is structured around the four stages of idea generation, task description, detailed expansion and task implementation. Each of these stages is further divided into work modules. For example, Task Description comprises modules on Task structure, Core activity and Inputs/outputs;
- Modifying an existing task; The user can select a task from the Examples database, and then modify it to suit specific classroom/examination needs;
- Task commentaries: These are equivalent to critiques, with the addition of facilities for documenting specific classroom/examination information. Task commentary is viewed as an integral part of all task development activities;
- Choosing a task: This module supports the selection of tasks and task sets from the examples database for fulfilling curriculum and assessment goals.

CONCLUSIONS

The studies reported in this paper show effects that contrast with standard expert/novice differences. The early reliance on depth-first decomposition by designers seems more like novice behaviour. Similarly, the delayed production of conceptually deep sorts contrasts with the standard finding that experts sort domain items at a greater conceptual depth. Of course, these are not standard expert/novice studies, since both teacher and designer groups have expertise. Instead, the studies are better characterised as studies of different specialisms, in which the kinds of conceptual representation and strategy appropriate for one are not necessarily optimal for the other.

Taking the results of Studies 1 and 2 together, they suggest that, while specialist task designers do possess deep conceptual representations of domain knowledge, they are able to suppress them when the constraints of an activity demand the search for original ideas. This is important in the light of demonstrations of impairment in expert performance resulting from the presence of deep conceptual knowledge (e.g., Wiley, 1998). Wiley has suggested that domain knowledge can act to inhibit creative problem-solving, though the task she used to explore expertise (Mednick's, 1962, remote associates task) was somewhat contrived and unnatural. Our results suggest that, when realistic activities that experts might expect to be engaged in are studied, creative problem-solving is not inhibited by domain knowledge. It may be questioned whether one can generalise across EFL and Mathematics domains, as in Studies 1 and 2. However, similar patterns of design practice to those found in Study 1 are evident in the data from an unpublished study of task design in the Mathematics domain we have undertaken as part of the ongoing project.

The results of the studies have a number of important implications for theories of problem-solving. They suggest that specialist designers are able to override the kinds of conceptual knowledge representation that might otherwise inhibit their creativity. Also, it appears that, even in highly creative domains, problem-solving is a highly structured process. The standard finding in the expert/novice literature regarding control strategies for problem decomposition is that experts typically work in breadth while novices work in depth (e.g., Jeffries et al, 1981; Adelson & Soloway, 1985; Davies, 1991). Our results demonstrate that, initially, specialist task designers work in depth while non-specialists work in breadth. Nonetheless, the specialists produced, in our view, some highly original and practicable solutions that successfully met the brief. Furthermore, the results show how problem-solving strategies are flexible across phases. Strategic shifts have been found previously in knowledge-lean problems, such as the missionaries and cannibals (Simon & Reed, 1976). The current research extends the study of strategy flexibility to a knowledge-rich, creative problem-solving domain. Changes in expert control strategy over time are a novel finding, and suggest that manipulation of control strategy is an important aspect of creative design expertise.

The outcomes of the study also provide us with a strong empirical basis on which to develop task design guides for non-specialist task developers. We developed the Task Wizard around a procedural model in which task development is guided by phases that encourages early depth-first exploration followed by later breadth-first control of expansion, with concurrent evaluation. Also, the data provide a rich set of task design criteria that are appropriate for evaluating emergent designs at different stages. Furthermore, we are able to encourage the search for contexts in which novel task ideas might be sought, and delay the provision of task exemplars until the ‘creative’ phase of idea generation has been achieved. At later stages, task exemplars and descriptions of task plans can be used as evaluative tools. We believe the studies demonstrate how an empirically-based cognitive science approach to studying expertise contributes to the design of effective technologies for supporting the design and adaptation of educational tasks.

ACKNOWLEDGMENTS

The research reported in this paper was supported by grants from the Economic and Social Research Council's (UK) Cognitive Engineering initiative, No. L127251031, and from the National Science Foundation ESI 9726403. We thank Catherine Fritz and Shona Saul who assisted in the data collection and analysis, and all the designers and teachers who took part in the study. An extended report of Study 2 was presented at the 21st Annual Conference of the Cognitive Science Society, Vancouver, August 1999 (Ormerod, Fritz & Ridgway, 1999).

REFERENCES

Adelson, B., & Soloway, E. (1985). The role of domain experience in software design. IEEE Transactions on Software Engineering, SE-11.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation in physics problems by experts and novices. Cognitive Science, 5, 121-152.

Davies, S. P. (1991). Characterizing the program design activity: Neither strictly top-down nor globally opportunistic. Behaviour & Information Technology, 10(3), 173-190.

Eason, K. (1988). Information technology and organizational change. London: Taylor & Francis.

Guindon, R. (1990). Designing the design process: Exploiting opportunistic thoughts. Human-Computer Interaction, 5, 305-344.

Jansson, D. G., & Smith, S. M. (1991). Design fixation. Design Studies, 12, 3-11.

Jeffries, R., Turner, A. A., Polson, P. G., & Atwood, M. E. (1981). The processes involved in designing software. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. Cognitive Science, 14, 511-550.

Mednick, S. (1962). The associative basis of the creative process. Psychological Review, 69, 200-232.

Ormerod, T. C., Fritz, C., & Ridgway, J. (1999). From deep to superficial categorization with increasing expertise. In M.Hahn (Ed.), 21st Annual Conference of the Cognitive Science Society, Vancouver: LEA.

Perez, R. S., Fleming Johnson, J., & Emery, C. D. (1995). Instructional design expertise: A cognitive model of design. Instructional Science, 23, 321-349.

Schoenfeld, A. H., & Herrmann, D. J. (1982). Problem perception and knowledge structure in expert and novice mathematical problem-solvers. Journal of Experimental Psychology: Learning, Memory & Cognition, 8, 484-494.

Simon, H. A., & Reed, S. K. (1976). Modelling strategy shifts in a problem solving task. Cognitive Psychology, 8, 86-97.

Soloway, E., & Erlich, K. (1984). Empirical studies of programming knowledge. IEEE Transactions on Software Engineering, SE-10, 595-609.

Visser, W. (1994). Organisation of design activities: Opportunistic, with hierarchical episodes. Interacting with Computers, 6(3), 235-274.

Wiley, J. (1998). Expertise as mental set: The effects of domain knowledge increative problem-solving. Memory & Cognition, 26, 716-730.

Wirth, N. (1971). Program Development by Stepwise Refinement. Communications of the ACM, 14, 221-226..