

Styles and Approaches to Problem-solving

DIANA LAURILLARD

Introduction

Problem-solving tasks are set as a regular part of the course work on most courses in science, mathematics and technology, and in some social science courses as well. They are seen as an important part of the students' work because they require the application of knowledge and principles to new situations, thus testing and reinforcing the students' real understanding of what they have learned. Knowledge without the ability to apply it is rightly seen as a very poor commodity, and teachers therefore regard problem-solving exercises as an important part of learning.

We can assume, for the purpose of this chapter, that the problems being set for students have a purely educational value; that it is not so much the solution that is of interest, as the process of reaching that solution. We can thus define a problem-solving task as one which 'engages the students in thinking about the subject matter in ways designed to improve their understanding of it'. Problems may sometimes be set to give students practice in some procedure, such as solving quadratic equations, but students learn little from this, other than a facility with the procedure itself. Such problems do not fit into our definition. We are concerned only with problems intended to develop in the students at least a greater familiarity with their subject, and perhaps a better understanding as well.

The teacher faces a difficult challenge in designing problem-solving tasks that fully serve this educational function. Such tasks should help the students to weave the factual knowledge they have into their own conceptual organisation, by enabling them to elaborate the relationships between concepts and to impose structure on the information they have. If they do less, then the exercise can easily become a meaningless mechanical manipulation, and loses its real educational potential. Naturally, for many teachers the choice of problem-solving tasks is circumscribed by the traditions of their subject, and there is relatively little creative effort involved in designing such tasks. Even when there is, it is more likely to be for the sake of the elegance of the problem, rather than for its educational value. But the design of problems is important because the cognitive activity inherent in a particular problem-solving task determines the way the student will think about the subject matter. 'Bookwork' problems will encourage bookwork solutions, requiring very little cognitive effort on the part of the student. A more imaginative problem that challenges the student and invites him to construct new ways of combining information will promote a better understanding. The point is illustrated neatly by Dahlgren's question to economics students about the cost of a bun (Chapter 2). They were practised at defining the laws of supply and demand, but

their lack of basic understanding was revealed by their inability to break out of familiar patterns of thinking to answer a very basic but unusual question.

If we can establish the characteristics of a good problem-solving task we must then ask how successful it is in practice. This brings us back to the main theme of this book. Here we ask "what are students' experiences of learning from problem-solving?"

In this chapter, we begin by considering how problem-solving has been studied in the past, and how this relates to recent studies of the students' experience of problem-solving. We shall find that students' approaches to problem-solving can be described in terms of the deep and surface approach already introduced in Chapter 3. This categorisation is developed further to include a theoretical analysis of the internal relations between the students' learning processes and the nature of the subject matter content. The aim overall is to clarify the nature of learning from problem-solving which may then enable us to use it more efficiently as a teaching method.

Ways of Approaching an Understanding of Problem-Solving

Human problem-solving has been a continuing concern of psychologists, and they have developed different ways of investigating it. In this section, two well-established approaches are introduced, namely *Gestalt* psychology and Human Information Processing, while the next section develops a critique of them based on empirical studies using qualitative methods.

There are important differences between these two theoretical analyses of problem-solving. Gestalt psychology describes human cognition in terms of the quality of our perception and thinking, while information processing theory categorises the *mechanism* of our perception and thinking. Not surprisingly, therefore, the two types of theory produce very different descriptions of problem-solving.

Gestalt Theory and Problem-Solving

The essence of Gestalt psychology is to emphasise the structural quality of the way in which we perceive, think about, and feel, the world around us. This structural quality is wholeness ('Gestalt' means 'whole'). In order to see something, we focus on some part of it – like a word on a page. We select a part from a whole. In focusing on the foreground or 'figure' we thereby create a background or 'ground'. The essence of our perception is that each part exists by virtue of its relation to a whole, and can itself be seen as a whole. By emphasising this structural quality of human cognition, the Gestalt psychologists make the assumption that there is always some underlying structure within our perception of a situation, or experience, or task. They also regard relationships between parts and wholes within that structure as constituting the forces that drive our productive thinking.

Wertheimer (1959) applied these ideas to exercises in elementary geometry to show how Gestalt theory can be useful in understanding problem solving. The theory suggests, for example, that the best way of discovering how to find the area of a parallelogram is not by being taught a rule or algorithm, but by finding

the underlying structure of the problem, and thereby solving the problem in a meaningful way. The reasoning process might run as follows: the parallelogram is essentially a rectangle in the middle, plus two extra triangles:



FIGURE 8.1(a)

We know how to find the area of a rectangle, so the area of the middle part is known. We are left with the two triangular parts. They are not rectangles, but by rearranging the diagram they do fit together:



FIGURE 8.1(b)

and that makes one large rectangle with the same area as the parallelogram. Hence the problem is solved as 'area = length x height', where the reason for this is now apparent: the solution is generated from the visual restructuring of the problem statement. Wertheimer (1959, p. 239) describes this kind of process as follows (my parentheses):

When one grasps a problem situation, its structural features and the requirements set up certain strains, stresses, tensions in the thinker. What happens in real thinking is that these strains and stresses (e.g. what to do about the triangular parts) are followed, yield vectors in the direction of improvement of the situation (i.e. they fit together to make a rectangle which it is easy to find the area of), and change it accordingly (i.e. draw the reconstruction). (The solution) is a state of affairs that is held together by inner forces as a good structure in which there is harmony in the mutual requirements (i.e. the reconstruction is equivalent in area to the original, but also allows us to calculate the area) and in which the parts are determined by the structure of the whole, as the whole is by the parts.

What Wertheimer has done here is to explain the process by which we can solve a problem, not in terms of a procedure, or a series of steps, or even a strategy, but in terms of the way in which we perceive the whole problem situation. The forces that drive our thinking along the steps to the solution are created by our perception of the structured requirements, in other words what the 'givens' of the

problem need to become the solution. Such an account of problem-solving emphasises the importance of the meaning of the problem for the student. When we draw on Gestalt theory to think about problem-solving, it is inconceivable to think of teaching children to solve problems by some rote method.

There are two main difficulties in applying Gestalt theory to the kinds of learning and problem-solving that occurs in the classroom. One is that the problems researched are of a particular character – geometric, algebraic, mathematical. It is not clear how far the theory can help us with different kinds of problems, (e.g. experimental situations or engineering problems) which have very different structural characteristics from those often discussed in the literature. The second problem is that the focus is always on the problem and the student's perception of it. But from the student's point of view, the problem situation is not just the content of the problem as given but includes also the context in which it is given. Wertheimer himself makes the same point in his introduction to *Productive Thinking*, p. 12.

The nature of the topics discussed permits us to deal with thought in terms of "relatively closed systems", as though thinking about a problem were a process that occurred independently of larger issues. Only occasionally shall we refer to the place, role and function of such a process within the personality structure of the subject and within the structure of his social field.

Again, in the conclusion, he describes the problem-solving process as:

. . . a partial field within the general process of knowledge and insight, within the context of a broad historical development, within the social situation, and also within the subject's personal life. (p. 240).

We can imagine this broader context by considering the problem from the student's viewpoint. Does he just have to find the area of the parallelogram, or does he also have to do it in the way the teacher wants? If so, he may wonder whether or not he could get away with doing it his own way, or even consider the consequences of not doing it at all. It is a far more complex 'problem' than we might at first suppose, and all these issues have some kind of bearing on what precisely the student does with the content of the given problem, as we shall see later in the chapter.

Information Processing and Problem-Solving

The Gestalt account of problem-solving tells us that the structural quality of our perception assists the solution process, and when we fail to solve problems, this amounts to a failure to perceive the structure of the problem situation. By contrast, information processing theory focuses on the mechanism of the problem-solving process. Both theories begin by looking at the ways in which people go through a problem-solving process, but they do it in different theoretical contexts, and so focus on different aspects of the situation. Information processing looks at the procedures that people adopt, and integrates these into a more deterministic account of how humans solve problems. It is characteristic of this type of account that it should be capable of supporting a computational model, which "aims at the

representation of a psychological theory of problem-solving". (Boden, 1978, p. 143). This approach to theory within cognitive psychology led to the development of a new field within instructional design, which experimented with the computational modelling of students' problem-solving processes, especially for the construction of intelligent tutoring systems. The origins of this approach can be traced to the work of Newell and Simon, who developed a program called 'General Problem-Solver' (Newell and Simon, 1972). They derived a psychological theory of human problem-solving from an analysis of people attempting to solve 'brain teaser' problems. The theory was based on the idea that human cognition is dominated by heuristic processes. Their analysis of protocols revealed these heuristics, which could then be represented in a computer program capable of solving the same problem in a similar way. Failures to solve problems could then be seen as failures either to apply the correct heuristic, or to use one at all.

The general heuristic procedures, such as means-end analysis, creating sub-goals, or working forwards and backwards can be applied to any problem. The General Problem-Solver used these heuristics, together with an appropriate representation of the problem, to generate the specific heuristics for that problem. This, the theory states, is what a human will do when confronted with a new problem, i.e. use general heuristic procedures, together with an appropriate representation of the problem, to generate a specific solution. The value of the theory lay in its description of the heuristics of human problem-solving in a form capable of computational modelling. This opened up the possibility that, if computers could model ideal human problem-solving, then they could also be expected to model flawed problem-solving procedures, by perturbing that ideal in specific ways. This would be valuable in an instructional context if, by modelling a flawed problem-solving strategy, the program could generate the same incorrect result as a student. Since the program would then have a representation of the flawed strategy (e.g. as deleting one step in the correct procedure), it would be able to advise the student on how to correct the flaw (e.g. "Have you forgotten the following step?"), and thereby provide individually adaptive tuition. This intriguing idea led to a number of computer-based experiments with attempts to model students' problem-solving behaviour, mainly in mathematical subjects (see Wenger, 1978, Chapters 9 - 12 for a survey of these experiments).

One of the greatest theoretical difficulties with the information processing approach is that it begs the very important question of what is an appropriate representation of the problem. Some of the research in the field of intelligent tutoring systems has attempted to answer this by analysing students' problem-solving procedures in comparison with expert approaches. By modelling the student's problem-solving procedure as a perturbation of the expert's, it is possible for a computer program to generate remedial teaching from the nature of the perturbation, e.g. the student can be reminded of the omission of a vital step, or if an incorrect rule has been inserted into the procedure, they can receive remedial teaching on that. However, as I have argued elsewhere, this kind of analysis locates the student error at an inappropriate level of description (Laurillard, 1988).

The particular omission or incorrect rule may often arise for the student from an underlying learning problem, such as a misconception or a misrepresentation of the structure of the problem. In such a case, remedying the resultant error will not remove the underlying problem. A program can model the procedural aspects of a student's approach to a problem, but not the conceptual aspects. If the mistake in a subtraction problem is to insert the rule ' $0 - n = n$ ', for example, the program may be able to model this, and hence diagnose the error at this level, but what is the conceptual representation that allowed the student to entertain such an idea in the first place? The program has no access to that, and yet that is where the underlying learning problem probably lies.

The Gestalt theorists set out to describe the underlying structure of a problem from the expert's point of view, which can be used within instruction to direct learners towards the most appropriate form of representation of the problem. Information processing theories described how experts proceed through a problem, and more recently, student modelling studies set out to compare how novices proceed through a problem. Both theoretical frameworks analyse approaches to problem-solving, and produce complementary findings – the one on the importance of perceiving the underlying problem structure in an appropriate way, the other on the importance of following the appropriate solution procedure. We might expect that we could combine the two to give a complete picture of how students solve problems. The question is: how far does the theory apply to the practice of problem-solving?

The Application of Theory in Practice

To achieve an understanding of the students' experience of learning through any medium, it is necessary to develop rather different methodological procedures from those described in the previous section. We are no longer concerned with describing the general processes of human psychology, nor with the computational modelling of procedures, but with the personal reality experienced by students as they learn. In order to understand learning from the learner's perspective, we need to use investigative methods that are capable of probing the students' learning experiences, and of eliciting data that will give us some insight into the full complexity of the learning process as practised by students.

A study of how students learn from problem-solving was carried out with a group of 12 students studying the second year of a combined science course at a British university. The aim of the study was to investigate how students approach and carry out problem-solving tasks set as part of their coursework. The course chosen was a course on micro-electronics, and the study focused on three of the problems set. For each problem, the student was asked to complete a short open-ended questionnaire, including such questions as 'how did you start the exercise?', 'were there any points you found difficult — what did you do about that?' and so on. Questionnaires were completed soon after the problem task had been finished so that students were able to remember what they did in some detail.

The questionnaire data were analysed by searching for students' descriptions of the kinds of heuristic activities defined by Newell and Simon. Several such

heuristic devices were apparent, but they did not operate in quite the way we might have expected.

Data of this kind necessarily give us a different perspective on the process of learning. They cannot tell us what cognitive processes are involved and how they operate, but instead they can tell us, for example, how the student perceives the given problem-solving task. Consider these quotes from students, explaining their initial approach to a problem which involved writing a device control program for a given microprocessor. The quotes record the important first step of making sure they understand the problem.

I read through the question to see what was familiar from the lecture, i.e. phrases or specific words that were repeated.

I have to sort through the wording very slowly to understand what he wants us to do.

I read through with reference to the class notes making sure I understand the sequence.

First I thought: the drawn circuit was incorrect: experimented with the current version (in the) notes.

The students' descriptions of their initial approach to the problem vary in the degree of activity involved; the first one is clearly taking a 'surface approach' (as described in Chapter 3), and the last is active to the point of being critical of the question. But what is common to all these students is the focus of their attention, not on the problem itself, but on the problem as set by a teacher in the context of a particular course. We might expect that the first stages in solving such a problem would be to consider what kind of microprocessor it is, what kinds of control would be needed, which instructions are relevant etc. But the students' attention is focused not on the program to be written, but rather on what they think the teacher requires. The domain of the students' problem is not the world of microprocessors but the world of the teacher setting this question. Each student, in different ways, relates the problem to its educational context: the lecture, the lecturer, the lecture notes.

The same tendency can be found in students' descriptions of how they carry out the problem task; again, they may be working not just on the content of the problem in isolation, but on the problem in context:

I thought of a diagram drawn in a lecture and immediately referred back to it. Then I decided which components were wanted and which were not and started to draw it out, more or less copying without really thinking.

I just looked back at the class example and tried to think of how it was similar and how it differed and where I would fit in the new stages of 'initialisation' and 'recovery'.

I decided since X was setting the question, block diagrams were needed.

Each step, and each strategic decision made, refer to the immediate context of the problem as it occurs in that course. The criterion is not "is that what this type of microcomputer needs?", but "is this what this teacher is looking for?". For these

students, the problem situation is quite different from those featured in experimental studies. The problem is not an isolated event; as Wertheimer said, it occurs "within the social situation" (op cit); it comes after a certain lecture and is likely to relate to it. It will also be marked by a particular lecturer, and the solution should take that into account as well.

The final stage of any problem-solving process involves checking back and examining the solution. And again, some students evaluate their solution not within the terms of the problem alone, but in relation to other aspects of its context, such as their own level of commitment to the task.

I don't think the finished product was right, but I decided it would do.

I drew what I thought seemed logical although [I] was not satisfied as I couldn't really see how it fitted in . . . I didn't really do this exercise with a view to getting anything out of it. I felt it was something to copy down and nothing to understand really.

I went back to re-check again the answer, not only to make sure that it was correct but to make sure that I had understood what I had written.

These illustrative quotes show that it is possible for the student to be so concerned to solve the problem *in its educational context* that they pay little attention to the problem itself or its inherent subject matter content. Teachers should always be aware that the student's perception of a problem is likely to be different from their own in this respect, and that this difference may have an undesirable effect on the problem-solving heuristics the students adopt. These heuristics, as illustrated above, are perfectly reasonable and reflect the rational procedures outlined by the information processing theorists, but if the content of the students' problem is this curious 'problem-in-context', then the content of their heuristic thinking will have less to do with microelectronics or whatever the subject matter is, and more to do with the requirements of their course. So the problem-solving task may fail to ensure that the student learns about the subject matter. They do learn something about the 'problem-in-context', but that includes knowing about how to get good coursework marks, about reading between the lines, about interpreting the lecturer's behaviour, and so on. This will no doubt be of some value to them, but it will not satisfy the teacher who intends them to learn about microelectronics. This is a serious difficulty in making sure that a problem-solving task evokes successful learning.

We have seen that students pay considerable attention to peripheral aspects of the problem, but students do also have to consider the problem itself. In the next section, we consider alternative ways of approaching a problem-solving task and attempt to determine the origins and consequences of these.

Students' Approaches to Problem-Solving

A more wide-ranging study of students' problem-solving activities was carried out on a group of 31 university students studying a range of science and engineering courses (Laurillard, 1978). The primary aim of this study was to find out how far existing descriptions of the learning process, such as the deep/surface approach, were applicable to students taking these courses. Each student was interviewed

on at least three occasions about a coursework problem task they had been set. The interview lasted one hour and included three stages:

<i>Teachback</i>	the student 'teaches' the problem situation to the interviewer, who interposes no substantive questions (based on the work of Pask, discussed later).
<i>Stimulated recall</i>	the students are interviewed about how they worked on the task in detail, using the problem statement and their written work to stimulate recall.
<i>Questions on context</i>	the students are interviewed about why they did what they did, relating this activity to other aspects of the learning context such as the relevant lecture, tutorial, assessment etc.

Transcripts of the interviews were analysed by searching for descriptions of, for example, a deep or surface level approach, and their interpretations were then independently checked by two other judges. Such data are rich not only in confirmation of the existing descriptions of learning, but also in insights into how the students experience these particular learning tasks. In Chapter 3 the students' approaches to reading were discussed in terms of the deep and surface approach, and it was demonstrated that the two forms of activity led to different learning outcomes for the students. It was also indicated there that these descriptions of the learning process do not apply only to reading. And here we find similar activities in students' approaches to problem-solving. The deep approach is active, with students looking for the meaning of what they are doing. The following quotes are taken from interview protocols of students describing their approach to various coursework problems in science and engineering. Each one illustrates the student's concern to focus on the meaning of the problem:

What I'm trying to do is picture what's going on and see the model they're using.

I was trying to work out what's happening to this point moving on a surface.

First I read the introduction to see what they had to say about it, why it gives a reasonable approximation and what it neglects, because you have to realise the limitations of the method.

Looking at the system, I was thinking of what is actually happening, relating numbers to features.

It is therefore possible to show that there are circumstances in which student's attention is focused fully on the subject matter content of the problem, and thus that the 'deep approach' describes one aspect of how students learn from problem-solving. Similarly, the alternative 'surface approach' can also be found. This is more passive, with the student content to treat the elements of the task in a purely mechanical way, not considering their meaning, merely their form:

I didn't really look at my notes because you don't have to actually look at the system, you don't have to interpret it in terms of its application.

I just copy from last year's notes . . .

You can't really go wrong, it's all done on the diagrams for you, you can go through without thinking at all.

The key to the deep/surface dichotomy, as found in reading tasks, is the focus of the student's attention; whether it is the meaning behind the words, or the words themselves. In applying the dichotomy to problem-solving tasks, we find an exact parallel, with students focusing either on the meaning, or on the words, numbers and diagrams themselves. The deep and surface approaches to learning can therefore be seen as characterising a fundamental aspect of how students learn, applicable in different types of learning task.

The origin of a student's approach to a particular problem is not apparent from their descriptions of how they work it out, but the interview questions about their perception of the educational context within which they are doing it, indicate that the approach derives from their intention — why they are doing it and what they expect to get out of it. Students who describe a deep approach in carrying out the tasks, respond to those questions with descriptions such as the following:

I want to understand the theory of what I'm doing to do a good write-up and get the results.

I have to use this for my project. I want to do as much of the steps as I can, to understand what's going on.

These quotes illustrate the students' intention to understand the meaning. In contrast, the surface approach derives from an intention merely to memorise or to reproduce:

These are general notes. It's an easy way of putting down principles so you can revise it.

I tend to write down certain things I rely on myself remembering for the next year or two ... you can remember it that way.

If the origin of the approach is the student's intention, then as the student may have different intentions within different learning situations, the same student may use either approach, on different occasions. In this study 19 out of the 31 students exhibited both types of approach (Laurillard, 1979). The internal consistency between intention and approach is illustrated by the following quotes from the same student talking about two different learning tasks:

Deep Approach

This has to be handed in — it's an operation research exercise, a program to find a minimum point on a curve. First I had to decide on the criteria of how to approach it, then drew a flow diagram, and checked through each stage. You have to think about it and understand it first. I used my

knowledge of O.R. design of starting with one point, testing it and judging the next move. I try to work through logically. Putting in diagrams helps you think clearly and follow through step by step. I chose this problem because it was more applied, more realistic. You can learn how to go about O.R. You get an idea of the different types of problem that exist from reading.

Surface Approach

This problem is not to be handed in, but it will be discussed in the lecture because the rest of the course depends on this kind of thing. I knew how I'd do it from looking at it; it practically tells you what equation to use. You just have to bash the numbers out. I knew how to do it before I started so I didn't get anything out of it. There's not really any thinking. You just need to know what you need to solve the problem. I read through the relevant notes, but not much because you don't need to look at the system. It's really just a case of knowing what's in the notes and choosing which block of notes to use. You don't have to interpret it in terms of the system. It's only when things go wrong, you have to think about it then. In this sort of situation you've got to get through to the answer.

Thus the deep/surface dichotomy does not characterise a stable characteristic of the student, but rather describes a relation between the student's perception of a task and his approach to it. The student's perception of a learning task encompasses a multitude of things: it depends on its form and content, on its relation to other tasks, on the student's previous experience, on the student's perception of the teacher who marked it and of how it will be assessed. But the operational outcome of this combination of judgements and perceptions is an intention either to understand or to memorise, and thereby to use either a deep or surface approach.

Thus the referential character of the deep/surface dichotomy – its description of what the student attends to – has been shown to be relevant to how students learn from problem-solving. The dichotomy has implications, however, for the way the student engages with the subject matter, and this is of crucial importance in problem-solving. This relational aspect of the dichotomy was described in Chapters 3 and 4 in terms of the distinction between 'holistic' and 'atomistic'. The terms define students' activities in carrying out the task. As we saw the 'holistic' approach involves students in attempts to "search for the author's intention to relate the message to a wider context and/or to identify the main parts of the author's argument and supporting facts" (Svensson, 1977). The 'atomistic' approach involves students in "focusing on specific comparisons", focusing on the parts of a text in a sequence (rather than the more important parts), memorising details and direct information indicating a lack of orientation towards the message as a whole" (Ibid.) The holistic/atomistic dichotomy focuses on the way students manipulate the structure of the text they are reading, and thus makes clear how the differences in outcome arise: the difference in approach constitutes a difference in outcome by virtue of the fact that the students are interacting with the subject matter in a different way.

The holistic/atomistic dichotomy is again mirrored in students' descriptions of their approach to problem-solving tasks. The parallel to the holistic approach is manifested when students describe ways of dealing with the problem content that preserve the structure and meaning of each part and its relation to the whole.

I started by (deciding) what I needed to prove. I tried to set up in my mind how I was going to do it.

You do it by putting things in boxes, forget what's inside them and look at the whole picture.

You're told so much, you need to find some kind of relationship.

Contrast these statements with those that illustrate an 'atomistic' approach which ignores the structure of the problem and concentrates on cobbling together a solution by manipulating the elements rather than understanding the whole.

First, you have to isolate what one knows, or what facts are known. Then, consider what expressions to use.

I started by writing down equations, but you should start by thinking of what you need.

I looked up the formulae and made calculations from those.

The essential difference between a holistic and atomistic approach is that whereas the former preserves the underlying structure of the subject matter content, the latter effectively distorts it, because the students pay no attention to the structure and concentrate only on juggling the elements together until they fashion a solution.

This structural aspect of approaches to problem-solving, which the holistic/atomistic dichotomy emphasises, is crucially important. The whole point of problem-solving as a learning task is that it should engage the students actively in thinking about the subject matter, and in operating on the relations within it, so that personal meaning can be created. The evidence from these interviews demonstrates that the two alternative approaches to problem-solving do exist, and clearly one is desirable and the other less so, at least if students are to be effective problem-solvers outside the narrow educational context. But we need a full understanding of how deep and holist approaches lead to a higher level of learning outcome if we are to make use of this finding in designing problem-solving tasks. What does it mean for a student to understand a topic, and how do different approaches to learning relate to understanding? The next section introduces a theoretical analysis of these questions from which we can derive a further way of describing how students learn from problem-solving.

Problem-Solving Tasks and their Relation to Understanding

The studies reported in this book deal mainly with studies of the learning process as seen from the student's point of view. The power of this type of research is that it allows us to investigate a process that is essentially internal by obtaining students'

descriptions of their experiences of learning. Such descriptions refer to the structural aspect of human cognition, identified by the Gestalt psychologists, but they are elaborated in relation to the particular context of higher education. Thus we find, for example, that structure can exist in both holistic and atomistic forms.

The problem of investigating the internal process of learning was solved in a different way by Gordon Pask and his colleagues (see Pask, 1976). They attempted to 'externalise' the process by creating an external manifestation of its most important features. One of the techniques adopted was to arrange the factual and descriptive information about a subject on a series of cards, each one labelled with a description of its content. Putting the information together would enable students to work out underlying principles (for example, of biological classification). Students were required to work out those principles (i.e. they had to demonstrate understanding) by selecting and reading the cards in any order they chose. The organisation of the subject was essentially hierarchical, but students could work in any way they liked, for example, from general points to specific ones, or vice versa, or across the topics. Pask found two contrasting strategies; for one, students built up the complete framework by beginning with general descriptions and filling in details later. For the other, they built up the framework step by step from the details to the more general principles. The two strategies achieved equivalent outcomes because that was required by the task. Thus there appeared to be two distinct strategies, ultimately equivalent in outcome, but very different in process.

This methodology is fundamentally different from that used by Svensson and Marton. They, like the Gestalt psychologists, make the students' perceptions of the structure of the material the focus of their investigation. Pask, like the information processing theorists, takes the structure of the material as given, and investigates what students do with that information, or how they process it. This has become the dominant mode of investigating problem-solving in educational contexts (see Wenger, *op cit*).

The result of Pask's investigation is interesting, but not immediately applicable to normal teaching-learning situations because the learning task was so artificial. However, in parallel with this experimental work, Pask also developed a theoretical framework, which allows us to interpret his findings and apply them to more familiar learning tasks.

Pask developed Conversation Theory as a way of describing the logical structure of what an individual (person or even machine) must be able to learn, and what the nature of the relation is between this individual and the subject matter to be learned. One basic principle of the theory is that in order to be able to learn something, an individual must know what it knows — must be able to tell itself what it knows, hence 'Conversation' theory. The second basic principle is that the individual must come to know a subject domain, must operate on it (manipulate its elements according to some plan or procedure) and must obtain feedback on the result of these operations. Finally, the third basic principle is that in order for these operations to form a systematic well-organised investigation of the subject matter domain, they must be generated from a global theoretical

framework through a set of operations and procedures which also receives feedback on the results of its operations. An obvious parallel is global theory generating localised experiments in scientific method — the prototypical way of learning about a domain. These basic elements of Conversation Theory (see Pask, 1975, for a fuller description) can be combined and represented as the diagram in Figure 8.2.

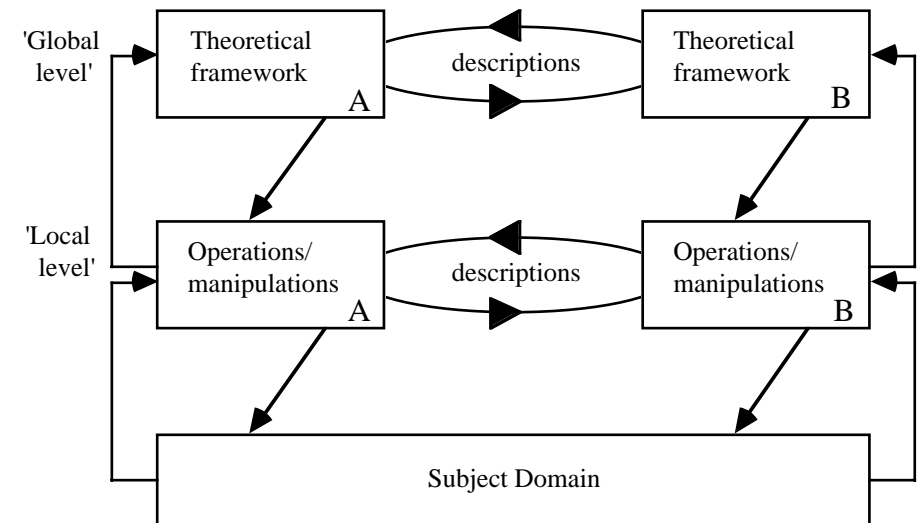


FIGURE 8.2

Schematic representation of Conversation Theory (based on Pask, 1976).

The symmetry of the diagram represents the internal 'conversation' that constitutes learning. The two theoretical frameworks we may call A and B, may be different from each other, but must be operationally equivalent. For example, the first may define a circle as a polygon with an infinite number of sides; the other may define it as a line whose locus maintains a constant distance from a fixed point. When both are used to generate an output through the manipulation of lines and points, they will both generate circles. The 'descriptions' then refer to the way students make the content meaningful to themselves. Having constructed the meaning (e.g. a description of theoretical framework A) it is then possible to use it to construct B, an alternative but compatible framework. That alternative must then be tested by generating corresponding operations on the subject matter domain and checking that they produce the same results as the previous framework. The vertical pathways in the diagram may also be used to construct frameworks, as in scientific method. The three levels in the diagram indicate different aspects of the subject matter. First there is the global theoretical framework — the structured elements and their relation to each other and to the whole, then the localised manipulative procedures — the specifics, isolated details

unrelated to the whole; and finally the domain itself — the external representation of the subject. As an example, take Pythagoras' theorem as the theoretical framework, techniques such as constructing a square on a line as the manipulative procedure, and geometric triangles as the domain.

From this purely theoretical account, Pask derived two styles of learning, both of which are necessary for understanding, i.e. the proper development of the theoretical framework. 'Operation learning' refers to the vertical pathways: the construction of hypotheses, the use of rules, techniques, procedures, the manipulation of entities in the subject matter domain. 'Comprehension learning' refers to the horizontal pathways: the description of the construction at both levels, global and local, the interpretation of their meaning, the search for analogies with other similar constructions.

These are theoretical descriptions of learning, but they may nonetheless be applicable to the reality of student learning. In the research study already described (Laurillard, 1978), the students' work on their coursework assignments was also used to investigate the applicability of operation and comprehension learning in this kind of learning situation. Ten of the students were interviewed about three of their assignments, each one a problem-solving task — in chemistry (reaction kinetics), crystallography (stereographic projection) and metallurgy (equilibrium diagrams). The students were interviewed individually, and at the start of the session they were asked to do a 'teachback' (Pask's term), i.e. to teach back to the researcher what they had learned about the subject matter of the problem-solving task. Each teachback lasted 5-10 minutes and was recorded and transcribed for later analysis. The analysis was done by inspection, looking for examples of statements that described either operation learning (statements of rules or procedures) or comprehension learning (descriptions of concepts or interpretation of operational constructions). The analysis was checked by two judges who achieved an average of 82 per cent agreement in assigning these categories.

Given this form of analysis, the presence of operation and comprehension learning within students' normal academic work can be demonstrated by selecting quotes from protocols which have been classified according to the two styles. The following quotes, where each student is describing how to work out a problem on an equilibrium diagram, illustrate how they employ the lower, localized level of operation learning in the form of standard techniques in working out the solution. An equilibrium diagram represents the structural phases that metal alloys go through as they cool. The students were asked to work out the sequence of phases for particular alloys using the diagram. This involves them also in interpreting shapes and sections of the graph.

If we're at a certain point, we can find out the proportions of the length of the line.

You work from this side of the graph, you get 12.5°A; as temperature is raised, solubility is increased.

Now, bring in that rule, anything between those two single phases, you've got a double phase, so that's a double phase.

Now then, you've got a straight line. Now then, another rule is that if you've got a straight line . . . that is a compound.

These students are clearly using operations — procedures and rules — but they are not operating at the level of the theoretical framework of equilibrium diagrams. The focus of their attention is on isolated details of the subject matter, and operations are carried out on the basis of selection from a standard repertory of techniques rather than by recourse to theory. Similarly, we can find evidence of comprehension learning at the lower level, where descriptions of concepts are local, and there is no attempt to integrate concepts or establish relations between them.

This is eutectoid reaction here. This is your a phase. This is a two-phase region, which is a mixture of the a and carbon compound.

. . . this line . . . is called the liquidus, and by liquidus, it means everything above it is liquid.

If you've got pure iron and you elevate its temperature you get structured changes with increasing temperature, at 90° you start off with the first structure you call.

These students are focusing on the meaning or interpretation of the diagrammatic representation, but they are not descriptions of a theoretical framework, rather they are descriptions of its detail in isolation. Quotes of this sort indicate the presence of comprehension learning, but only at the lower, localised level. Evidence of learning at the more theoretical level can be found but in this study it was rare. One such example is still a description of structural changes, but here the meaning of the diagram is related to the theoretical concept of the crystal pattern. It is thus not simply a description of the existence of the phases as areas on the diagram, it places that interpretation in its theoretical context.

In some types of material, a lot of them when they freeze, metals that is, you get two distinct crystal patterns. In a particular metal you could end up with one phase with dendrites in it . . . they're two completely different phases and so, because a lot of metals aren't completely soluble when they start to freeze, you get these two phases out.

This student is offering an explanation of the theory to support his identification of the two phases of the diagram, and this is a form of high level comprehension learning, i.e. the student is building descriptions of the theoretical framework underlying the problem. Thirty such protocols were analysed in this way. All students were found to use both styles of learning, but in varying proportions and, more strikingly, in varying proportions depending on the task. For example, on the stereographic projection task, all the students showed a high incidence of operation learning, whereas on the Equilibrium Diagram task, only half the students did so, with half biased more towards comprehension learning. This unequal distribution of styles among different problem-solving tasks is strongly indicative of a task effect on choice of learning style, and this will be discussed further in the next section.

This research had thus demonstrated that the theoretical constructs of *operation learning* and *comprehension learning* also help to describe problem-solving tasks in everyday studying. An obvious question is, how are these constructs related to the descriptions of ‘approaches to learning’ we have already encountered?

We can begin to make sense of the relations between these constructs if we consider again their definitions. Operation learning concerns the manipulation of the concepts and objects in the subject matter domain. Comprehension learning concerns their meaning, or description. The global level involves integration of the descriptions into a theoretical framework: the local level does not. The descriptions of deep/holistic or surface/atomistic approaches do not involve a separation into procedures and descriptions. Thus the only parallel that can be drawn between the two sets of categories would suggest at least a tentative correspondence between deep/holistic approaches and both comprehension and operation learning at the global level, and between surface/atomistic approaches and both comprehension and operation learning at the local level.

If we consider approach, with its intentional component, as a preliminary to style, it is then possible to suggest that the choice of approach affords the opportunity for one or other level of style to be implemented. For any particular problem, a student who is thinking deeply and holistically will be looking for meaning and will be able to attend to the global level of descriptions, whereas the student who is thinking atomistically will consider only the local components of the problem without seeking to integrate them meaningfully. The effects of a surface approach, insofar as it involves the intention to reproduce, will be to produce low-level descriptions or unintegrated sets of operations.

A deep approach may go through the initial stages of low-level operation learning, but only as a preliminary to the high-level integration of descriptions and operations into a full understanding of the subject matter domain.

What Pask’s theory tells us is that for any problem, there are global and localised forms of description of its domain, and the student has to be able both to manipulate the concepts and the relations between them and to interpret the meaning of those manipulations. What Svensson and Marton tell us is that the global forms of description will not be considered by those students who take a surface/atomist approach, and they will achieve a full understanding of the problem only if they take a deep/holist approach.

It has been possible to show that the two forms of descriptions of learning, the one derived empirically, the other theoretically, are applicable to a wide range of normal academic tasks, and are compatible with each other. These descriptions give us a way of simplifying the complexity of students’ experiences of learning from problem-solving so that the task of trying to understand how students deal with this form of learning becomes more manageable. But how does this help us to use problem-solving more effectively as a form of learning?

Implications for Teachers

The student’s choice of operation or comprehension learning may depend as much on the nature of the task as on the student’s own personal characteristics. Some

tasks necessitate operation learning e.g. the stereographic projection problem required students to do considerable manipulation of mathematical objects, but did not require them to do any interpretation of the objects or the manipulations. Similarly, the Equilibrium Diagram task required students to interpret a diagram to give an account of what was happening to a cooling metal alloy, and this required some manipulation of objects and concepts as well. The empirical results confirmed that the requirements of the task, in each case, matched the predominant style of learning exhibited in students’ protocols (Laurillard, 1979). A similar result was reported by Taylor (1990), from an observational study of students solving problems in programming Prolog. Those teachers who took the declarative approach to teaching Prolog, which emphasises the logical rather than the computational representation of a problem, created a task that led students to embark on:

a simple surface ‘translation’ exercise when writing programs, rather than a transformation of the problem statement. (Taylor, 1990, p. 307).

Here again, the teaching essentially *required* students to focus on a logical representation of the problem. This emphasis to a failure to focus on the computational structure of the problem which they had to do if they were to construct a satisfactory solution. The process of ‘translation’ using operation learning was adopted, whereas a reinterpretation or ‘transformation’ using comprehension learning was actually necessary for success.

The choice of learning style has also been related to the student’s intention, as characterized by his approach to the task. But we must take care with the deductions we make here, because the categories of ‘approach’ have been derived from a reading task. There is an important difference between the two: a reading task does not itself make demands on the student — the text is there to be read as the student chooses, with some purpose in mind certainly, but the text itself does not state the purpose. A problem-solving task, on the other hand, explicitly requires the student to solve it. As in reading tasks, the student may approach the task with an intention to learn meaningfully or superficially and may choose how he carries it out, but the crucial point about a problem-solving task is that it may itself make very minimal demands. For many such tasks, there is a standard procedure which students are wise to adopt, but which need not engage them in thinking about the subject at a deeper level. Few such tasks really deserve the name ‘problem-solving’, and it is hard to find examples of genuine problem-solving in many degree courses. The students’ comments on their problem-solving strategies reveal how minimal some of these task demands can be. Thus the choice of approach may not derive wholly from the student’s intention.

We can see, therefore, that the student’s choice of approach and style is dependent to some extent on the nature of the problem-solving task itself, and also on how the requirements are perceived. Both of these influences are in the control of the teacher. If teachers wish the tasks set to be effective in improving students’ understanding of the subject, if they are meant to be more than purely mechanical exercises in rehearsing some standard procedures, then the design of those tasks is crucial. They must be complex enough to demand hypothesis-

testing or explanations of theory. The design process must take into account the various descriptions of learning we have discussed, and ensure that the problem requires the student to engage in the appropriate kind of thinking. It must also be considered in relation to assessment procedures and the whole educational context, as we shall see later in Chapter 13. After that, the responsibility for learning lies with the student.

Students take a largely rational approach to learning. They consider what is required of them, they decide on priorities, and they act accordingly. The teacher plays an important part in forming their perceptions of what is required and what is important, and it is this, as much as their style of presenting the subject matter, which influences what and how their students learn.