1 Peer interaction and learning: perspectives and starting points

Introduction

This chapter will set the scene for those that follow by going back to the 1970s to examine some of the strands of theory and empirical research which converged around the question of when and how peer interaction can facilitate children's understanding and learning. Starting with social learning theory, we shall develop a focus on the concept of socio-cognitive conflict as an engine of mental development, a concept that owes its origins to Piaget, via Doise and colleagues. The Piagetian origins of this idea are reflected in the selection of tasks in these early studies, and this in turn gives rise to certain problems in terms of the interpretation of some of the experimental findings. This consideration will take us on a slight detour in the course of this chapter, raising issues to be returned to later. The latter part of the chapter will be given over to an account of the series of experimental studies of peer facilitation of children's problem solving which marked our own initial engagement with this field of research.

Modelling success: the social learning theory approach

As inheritors of the behaviourist approach to learning, social learning theories exerted a strong influence upon the psychology of child development in the 1970s. Such theories saw *modelling* as a key formative process in cognitive as well as other aspects of development. Thus any facilitative effects of child–child interaction in learning were construed largely in terms of processes of imitation or modelling.

Studies were conducted in which children's performance on various cognitive tasks was assessed before and after they had watched other children performing the same tasks. Many of the tasks in question were drawn from the Piagetian repertoire, because social learning theorists were setting themselves against Piagetian constructivist explanations which they saw as unduly individualistic in emphasis (e.g. Murray, 1974; Rosenthal and Zimmerman, 1972). Such studies did provide some evidence that children who initially performed at 'pre-operational' levels on tasks such as conservation of quantity could be induced to give operational judgements on such tasks simply by being required to observe another child who offered such judgements. From a Piagetian point of view, however, such demonstrations were unconvincing since they did not offer evidence that the children concerned were able either to justify or to generalise these 'operational' judgements.

Other studies within this tradition went beyond passive observation to examine the effects of actual interaction between 'pre-operational' and 'operational' children in the context of such tasks (e.g. Miller and Brownell, 1975; Silverman and Geiringer, 1973). Here again, it was possible to show that the pre-operational children did tend to make progress, and in this case with some evidence that the children concerned could produce their own justifications for their new operational judgements.

However, in the interaction studies, it was harder to tie such progress to the child's exposure to a partner who 'modelled' the correct conclusion. A different kind of interpretation was developed in relation to essentially similar studies by researchers working within a 'constructivist' tradition.

Construction of understanding through socio-cognitive conflict

Piaget's own early writings (most notably, Piaget, 1932) offered an argument for the potential productivity of peer interaction in relation to cognitive development, and especially in relation to the achievement of concrete operational modes of thought in the early school years. Piaget saw the pre-school child's egocentrism as presenting the major obstacle to achievement of operational thinking. Such thinking required 'decentration', the ability to take account of multiple points of view, and more generally, multiple covarying factors in a given situation. Pre-schoolers tended to fix on the first relevant factor they identified, and to answer entirely in terms of that. What the child needed in order to progress was something which disturbed this 'centration'. Exposure to someone else who saw things differently, in a situation which called for resolution of the conflicting responses, was seen as providing just this kind of disturbance.

Importantly, Piaget considered that inequalities of power and status were inimical to the effectiveness of this process. If children were exposed to the response of a powerful figure such as an adult, they would be unlikely to take issue with it. Rather, they would tend to ignore it if possible, and comply with it if not. In the case of exposure to a differing point of view from an equal, however, the social dynamics of the situation would create a pressure towards resolution of differences: 'Criticism is born of discussion and discussion is only possible amongst equals' (Piaget, 1932, p. 409). Even if the second child's answer was as wrong as that of the first, the attempt to resolve their partial and 'centred' solutions would be likely to result in the achievement of a higher level, more decentred representation which could embrace what was correct in both of the initial offerings.

In the mid-1970s, Doise and colleagues in Geneva conducted a series of experiments on the effects of peer interaction on the transition to operational modes of thinking in five to seven-year-olds (Doise, Mugny, and Perret-Clermont, 1975, 1976; Mugny and Doise, 1978). These studies used a variety of Piagetian 'concrete operations' tasks, a favourite being a 'village' task based loosely on Piaget's famous 'three mountains' task (Piaget and Inhelder, 1956). Here, model buildings were arranged on a baseboard to form a little village. The buildings are oriented in relation to a fixed mark on the baseboard, depicting, say, the village pond. The whole arrangement sat on a tabletop in front of the child. To the side of the child was another table, with an identical baseboard, but perhaps oriented differently in relation to the child. The child's task was to use a replica set of model buildings to recreate exactly the *same* village on this second table. The task is more or less difficult depending upon the relative orientations of the two baseboards. Where a rotation relative to self is involved, pre-operational children will typically fail to take account of the reorientations necessary to preserve the relationships between the buildings and the fixed mark.

Children were first tested individually on the task, to get a 'pre-test' measure of performance. In a second session, perhaps a week later, children were given another opportunity to do similar tasks, but this time some of them worked in pairs or small groups while others worked alone. Assignment to conditions was essentially random, although in some studies allocations to particular groupings was on the basis of pre-test performance. Sometimes, when the 'village' task was used, the children were put in different positions relative to the tables, so that the necessary transformation was easier for one child than for the other. In a third session, a week later again, all of the children were given a post-test individually.

Through such studies, Doise and his colleagues were able to show that children of slightly different pre-test levels, working together in dyads or triads, tended to perform at a higher level when working as a group than children working alone. More importantly, this benefit carried over to the children's individual post-test performances. In other words, the extent of pre- to post-test progress was significantly greater for those children who had worked in pairs in the second session than for those who had worked alone.

Large differences between the children in terms of their pre-test levels were associated with less progress than small differences in initial ability (Doise and Mugny, 1984). Even children who showed identical levels of pre-test performance could benefit from working together if steps were taken to ensure that they would come up with conflicting responses. Thus, with the village task, Doise and colleagues arranged for such children to occupy different spatial positions relative to the array. This meant that their 'egocentric' responses would ensure that they came into conflict about where to place the buildings, even though they were both reasoning in the same way. Children paired under these conditions made more gains than similar children working on the tasks on their own.

Doise and colleagues interpreted their findings in terms of socio-cognitive conflict. The children in the pair or small group conditions found themselves confronted with solutions which conflicted with their own. This conflict, and the socially engendered need to resolve it, prompts the children to re-examine their own initial responses, and may lead them to recognise a higher order solution to the problem which resolves the apparent conflict (Mugny, Perret-Clermont and Doise, 1981). For this to occur, it is necessary that the children's initial solutions differ, but it is not necessary for any of them to be more advanced than the others, or for any of them to be correct. The real 'ratchet' driving the process is that resolution of children's partial or centred solutions can in the end only be found by adopting a higher level, more decentred solution, thus ensuring cognitive progress.

Perret-Clermont (1980) used essentially the same three-stage experimental design in a series of experiments on peer facilitation of conservation judgements. Pre-tests were individual, and involved a range of standard Piagetian assessments of children's understanding of conservation of quantity. Various arrangements were tried out for the children assigned to the social interaction condition in the second session. One which proved effective was to assign two 'conservers' and one 'non-conserver' to work together. The non-conserver was then given a task such as sharing out between the children by pouring from a jug into three different shaped glasses. The session went on until all the children agreed that they each had the same quantity. The non-conserving children exposed to this kind of interaction went on to show significant pre- to post-test gains on standard tests of conservation of liquid quantity.

With conservation, as with the village task, progress could result even from interaction between two non-conservers, provided that they generated differing initial judgements. These studies further highlighted the potential benefits of peer interaction for the development of children's thinking, and the supposedly key role of socio-cognitive conflict in underpinning them. They certainly helped to stimulate research interest in this area, and indeed triggered our own initial studies reported later in this chapter. Not surprisingly, however, they also engendered some controversy.

Reservations about socio-cognitive conflict

Doise and colleagues attracted a good deal of attention in the early 1980s, and not a little criticism. Russell (1981, 1982) argued that in tasks such as conservation the pre-operational child was in effect responding with an *opinion* of how the array looked after transformation. By contrast, the operational child was responding in terms of what is objectively the case. The difference, as Russell sees it, is one of 'propositional attitude'. Pre-operational children coming up with differing answers could quite well (and, Russell suggests, often do) simply agree to disagree. If and when conflict is effective, on this argument, it should be by prompting the children towards the adoption of an appropriate objective attitude, allowing them to bring to bear understandings that they already possess.

The idea that young children might in fact understand more about conservation than their responses on standard Piagetian tests would suggest was one of the main themes of Donaldson's influential book *Children's Minds* (1978). McGarrigle and Donaldson, as early as 1975, were able to show that manipulations to the context of presentation of conservation tasks, such as having the transformation of materials appear accidental, could have a dramatic effect upon children's responses. Donaldson suggested that the standard conservation procedure contained misleading socio-communicative cues to the child to attend to appearances, rather than to the actual quantities concerned. The very deliberate way in which the transformation of apparent quantity was effected made this the natural focus for the child's attention. When the transformation occurred as a seemingly accidental consequence of the activities of an errant teddy bear, the child was able to discount the transformation as irrelevant to the actual quantities involved.

Other similar studies followed. In some studies of our own, we were able to show that even where the transformation of materials was deliberately done by the experimenter, children would discount it if some plausible rationale for it was provided. Thus Light, Buckingham and Robbins (1979) used a badly chipped beaker as a reason for pouring the contents from one container to another, which just 'happened' to be of a different shape. This was done in the context of setting up a game for the children, who were tested in pairs. The game was such that it was important for quantities to be equal, but the chip in the beaker would make the game dangerous; thus the need to find another container. A substantial majority of five and six-year-olds judged that the resulting transformation did *not* affect the quantities involved, whereas almost all children of this age fail on the standard version of the same task.

This result has proved replicable (e.g. Miller, 1982), but the interpretation put on it at the time may well be the wrong one. That interpretation stressed the fact that the transformation was made to seem incidental to the proceedings, rather than central to them. Later research using different designs (Roazzi and Bryant, 1992) has failed to find significant effects for such 'incidentality'. On the other hand, our own subsequent studies have shown that setting the conservation task in the context of a *game* is sufficient by itself to improve children's performance.

For example, Light, Gorsuch and Newman (1987) presented pairs of five and six-year-olds with a heap of dried peas for them to divide into two equal heaps. These were then put into two rather differently shaped containers, and the children were asked whether the quantities remained equal. For some of them, all this was done in the context of setting up a game in which they were going to compete with one another to move their peas as fast as they could to a target container, using a straw. The children who encountered the conservation questions in this context were much more likely to respond correctly than those who encountered the same transformation and the same questions without the game setting. So it seems that when children are working in pairs and anticipating a competitive game, they construe the conservation test procedures differently. In this situation they remain resolutely attentive to the quantities involved, and are not readily distracted by the appearances.

Things are actually rather more complicated. It can be argued (Light, 1986) that in these modified versions of the conservation test the children are really just complying with the apparent wishes and expectations of the experimenter. Just as the standard versions of these tasks might lead children toward the wrong response, these versions may in various ways cue the correct response, even in children with no firm grasp of conservation. Indeed we have been able to show that similar modifications can elicit 'conserving' responses even in situations where the quantity in question is not in fact conserved (Light and Gilmour, 1983).

For present purposes, however, this is not really important. What is important is that serious question marks were appearing about the validity of these kinds of Piagetian tests, and the stability of children's responses to them. More particularly, there are clear suggestions from this literature that children in pairs or small groups may well interpret given tasks and questions differently from the way in which the same tasks and questions would be understood by individuals. Improvements in performance, at the time and subsequently, may reflect this altered understanding of the questions at least as much as it reflects any socio-cognitive conflict arising from different points of view within the group.

One thing that makes this alternative point of view attractive is that peer facilitation processes sometimes seem to work just *too* well. Thus, for example, Perret-Clermont (1980) found that social class differences in children's pre-test performance on conservation tasks were typically large, but after a session of interaction of the kind described earlier, the differences according to class disappeared. Similar findings have been reported for urban–rural differences, and ethnic differences (Light and Perret-Clermont, 1991).

If the pre-test differences were genuinely reflective of differences in the children's levels of achievement in this crucial area of cognitive development, is it really conceivable that such differences could be wiped out by a single session of ten to fifteen minutes of interaction? It seems much more plausible that the initial differences reflect differences in children's understanding of the meaning and reference of the conservation questions, and indeed there is some independent evidence that this is the case (Grossen, 1988).

It would seem that what disambiguates the questions for the children is the experience of sharing in the paired or group session, leading Light and Perret-Clermont to conclude that: 'the efficacy of the peer interaction procedure arises not (or not only) from the socio-cognitive conflict mechanism . . . but from the introduction of a norm of equality which serves to support correct responses, which are then carried over to the individual post-test' (1991, p.145).

This issue of how social norms influence cognitive functioning surfaced in a number of areas of research in the 1980s, not least in Doise's own work on the influence of 'social marking'. However, we shall leave the further exploration of this issue for a later chapter. In the remaining sections of this chapter, we shall turn from Piaget to Vygotsky to find a rather different set of 'starting points' for research on social (and more particularly peer) processes in learning.

Learning as the co-construction of understanding

Vygotsky's writings of the 1920s and 1930s, though they led to a robust tradition of research in the Soviet Union, had little real impact on Western developmental psychology for nearly half a century. In the 1980s, however, there was a rush of translation, commentary and exploitation of

'Vygotskian' approaches in relation to a whole range of research problems (e.g. Bruner, 1985; Wertsch, 1985; Newman, Griffin and Cole, 1989; Forman, Minick and Stone, 1993).

Vygotsky's work has perhaps contributed to this field in two main ways. The first rests on his attempts to characterise the fine-grained interpersonal interactions that take place in learning settings, and involves concepts such as the 'zone of proximal development' and 'scaffolding'. The second rests on his broader attempt to develop a 'cultural psychology', within which learning is seen to depend upon mediation by social, cultural and institutional processes at many levels. We shall explore both of these contributions.

Vygotsky was interested in the origins of what he termed the higher mental functions: thinking, reasoning and understanding. The development of these higher mental functions in humans was seen as a fundamentally social rather than individual process. The child's interactions with other people serve to mediate between the child and the world-to-be-learned-about, and so understanding learning depends upon understanding the particular types of interactions which serve to foster it.

The concept of a zone of proximal development is central to this approach. Children, or indeed adults, can be characterised in terms of what they can achieve unaided. Indeed most forms of assessment involve testing what individuals can do without help. But individuals may also differ in terms of what they can do *with* help. The attainments which are possible for an individual given a measure of support and guidance are, as Vygotsky put it, within that individual's zone of proximal development (ZPD). They are attainments that will be possible for that individual unaided at some point in the near future.

The concept of a ZPD is thus an integral part of a theory of teaching and learning. Tharp and Gallimore (1988) have elaborated this aspect of Vygotsky's work into what they call a theory of teaching as assisted performance. They see learning as a process of guided re-invention, whereby social guidance makes it possible for the learner to achieve a constructive intellectual 're-invention' of some piece of culturally elaborated knowledge. The emphasis upon understanding being a matter of construction is clearly shared with Piagetian approaches. The distinctive features are (i) recognition that much of what the learner needs to learn is already in some sense 'available' in the culture, and (ii) recognition that interpersonal processes play a key role in making that culturally elaborated learning available to the individual.

Not unnaturally, given this emphasis on guidance, Vygotsky saw the relevant interpersonal interactions as going on between the learner and a more capable 'other'. Indeed, he defined the ZPD as: 'The distance between the actual developmental level as determined by individual problem solving and the level of potential development as determined through problem solving under adult guidance *or in collaboration with more capable peers*' (Vygotsky, 1978, p. 86, our emphasis). As with Piaget, then, peer interaction is flagged as having a potentially important role in learning and development. But whereas Piaget's emphasis was on the status-symmetry of such interactions, Vygotsky's emphasis is more on the competence-*a*symmetry that will often be a feature of peer relations. As Tharp and Gallimore put it, to the extent that peers can assist performance, learning will occur through their assistance.

The concept which has been most widely used to capture the forms of guidance which support learners in their progress through the ZPD is that of 'scaffolding'. Introduced by Wood, Bruner and Ross (1976), it captures the sense in which, through encouragement, focusing, demonstrations, reminders and suggestions, a learner can be supported in mastering a task or achieving an understanding. To take the building analogy further, if we imagine building an arch with bricks it is easy to see the vital role played by the wooden 'formwork' used to assemble the arch. However, the role of this scaffolding is strictly temporary; when complete the arch will hold itself up, though without scaffolding it could not have been built.

Tharp and Gallimore see progression through the ZPD in terms of four stages. In the first, performance is directly assisted by more capable others through 'scaffolding' of one kind or another. In the second, the learner effectively takes over the role of the 'scaffolder' in relation to his or her own learning. This often means 'talking oneself through' a task, remembering requests, reminders and injunctions previously given, and so on. The third stage is marked by the falling away of such 'self-guidance', as performance becomes automatic. The fourth 'stage' just recognises the fact that we can get thrown back to earlier stages of the acquisition process by such stressors as tiredness, or by changes in the precise conditions of the task. Learning to drive provides a useful case in point for all of these stages of learning.

It is as a model of adult guidance of children's learning that most direct use has been made of these concepts. For example, Wood and colleagues (see Wood, 1986) conducted a series of investigations of how four-year-old children can be taught to assemble a 3D puzzle involving wooden blocks and pegs. First they observed mothers' attempts to teach their own children how to complete the puzzle. The mothers who succeeded best were those who shifted their levels of intervention flexibly according to how well the child was doing. This 'contingency strategy' can be seen as a way for the mother to gauge and monitor the child's ZPD as learning proceeds, and to provide scaffolding at just the right point.

In further studies Wood and colleagues showed that the adoption of a 'contingent' strategy by specially trained tutors also resulted in better learning outcomes than any of the alternatives explored. In their recent work (Wood and Wood, 1996, in press), they have been working towards the development of computer-based tutoring systems which will provide optimally contingent patterns of tutorial support; something which human tutors, even given training, find it extremely difficult to do.

Although Wood and colleagues are intent on improving upon 'what comes naturally', much of the research done from a Vygotskian standpoint has tended to see effective teaching and learning exchanges as essentially incidental to ongoing joint engagement in activities, whether between mothers and children in the home (e.g. Rogoff and Gardner, 1984; Rogoff, 1990) or between experts and novices in traditional craft practices such as tailoring or weaving (Greenfield and Lave, 1982; Lave and Wenger, 1991).

Attempts have been made, however, to use a Vygotskian approach to illuminate what is going on in the classroom. Perhaps the most successful is that of Edwards and Mercer (1987), who explored how, by skillfully guiding classroom discussion, the class teacher can establish and maintain a focus of shared attention, provide children with a language in which to describe their own experiences and, using that language, build up a body of 'common knowledge' about the topic in hand. On the other hand, attempts to apply the concept of scaffolding more directly in relation to classroom practice have in at least some cases (e.g. Bliss, Askew and Macrae 1996) been unable to find any evidence for such processes at all.

Vygotsky's ideas about the social-interactional bases of learning have also inspired a considerable number of studies of learning in the context of peer interaction. As one might expect, a good proportion of these involve interactions in which an older, more experienced or pre-trained individual is designated as tutor to a younger, less experienced or untrained peer. This kind of 'peer tutoring' has been shown to be effective both in experimental studies (e.g. Phelps and Damon, 1989) and in applied educational contexts (e.g. Topping, 1994). It may indeed have benefits for the peer tutors as well as the peer tutees (Barron and Foot, 1991).

As Hogan and Tudge (in press) note, there has been relatively little research on peer collaboration, as opposed to peer tutoring, approached from a Vygotskian standpoint. However, increasingly since the mid 1980s, Vygotskian ideas have come to colour the interpretations offered by researchers for peer facilitation effects observed in experimental studies.