

Constructivism, other theories of the teaching and learning process, and their relationships

Stephen W. Draper
 Department of Psychology
 University of Glasgow
 Glasgow G12 8QQ UK
 email: steve@psy.gla.ac.uk

Preface

This was a discussion paper for Ray McAleese's NATO advanced studies institute, August 1994., Heriot Watt.

Contents

Sections with numbers are independent principles

Part A: Basics. The philosophical nature of learning

- Introduction
- 1. Structuralism
- 2. Constructivism
- 3. Activity-orientedness (economising on learning)
- Piaget
- 4. Forgetting and reconstruction: multiple links

Part B: The content of learning (in HE)

- 5. Concepts vs. tasks as the content of learning
- 6. Learning descriptions: The socially distributed nature of knowledge
- 7. Schooling: special type of learning; power of surface structures; learning as an effortful task.

Part C: Relationships between learner and material

- 8. Motivation [prior & external. material, description, learning tasks]
- Deep and shallow learning
- 9. Selection of an immediate task

Part D: Learning and doing

- 10. Learning as a task
- 11. Learning through doing
- 12. Feedback
- 13. Teaching by requiring activity
- Laurillard
- Teaching techniques

Introduction

We need a theory of the teaching and learning process (TLP). We need it to design teaching, to make meaningful measurements of it, to understand it. It is important because teachers cannot cause learning in any direct and reliable way: it is clear that other factors are crucial, are having big effects on the outcomes. In fact often the book or courseware is only a small factor in what happens.

I have therefore been looking at the literature for theories of the TLP as it is in higher education. This paper is based on a look at some of the major landmarks: Piaget (of new relevance because of recent talk of constructivism), the work of Ference Marton and others on deep and shallow learning, and Laurillard's (1993) book in which she offers a model she derives from Gordon Pask. At this workshop, various teaching tactics or principles have been mentioned. However it is hard to find discussions of the relationship between these ideas: are they rival theories, or compatible pieces of the puzzle all of which must be taken into account in any adequate theory of the TLP?

In this paper I attempt to answer this, describing each topic briefly and where necessary separating them into component propositions which could be held or denied independently. (The numbered sections identify these separate propositions, whereas other sections summarise the work of individuals or discuss common claims in

the light of the elemental propositions.) To a great extent they are logically independent, even though often taken together. For instance Piaget's views can be summarised as structuralism, plus constructivism, plus activity orientedness: but it would be quite possible to hold any subset of these and not believe the rest. In some ways activity orientedness is pivotal: the view that learning is largely determined by the tasks that learners actively orient to. On the one hand, many constructivists including Piaget seem to believe that activity orientedness necessarily follows from constructivism, although I do not believe this: it seems to me to entail an additional plausible, but questionable, assumption. On the other hand, the most important ideas and work in (higher) education in recent decades can be seen as related in different ways to the interaction of activity with learning.

1. Structuralism

Expressed very crudely, in the context of learning structuralism is the view that knowledge comes in lumps, not as an amorphous heap of associative links. Cognitivists take this for granted, although if you consider a medical student learning the names of bones, or a chemist the names of the elements, then it is not so obviously relevant. It should be noted that not only B.F. Skinner but also connectionist theories are non-structuralist.

The second feature of structuralism is to contrast general declarative features of knowledge with implementation issues. Thus structure is seen as more important than issues of representation. Chomsky's competence models are a linguistic expression of a structuralist position, in which formal properties are more important than how they are represented or computed. It is the formal, structure that is seen as the important invariant, while computations and implementations are approximations to it that may vary idiosyncratically but unimportantly. Psychological models of resource limitations (e.g. working memory capacity), in contrast, are non-structuralist in nature.

The third feature of structuralism concerns the type of generalisation focussed on, and here the contrast is with functionalism. The latter is concerned with how mental functions (goals) can be decomposed into component functions (subgoals), and with how alternative methods for one goal might be compared and selected between. Structuralism in contrast focusses on what is common between classes of action or types of object, without primarily organising these in terms of functions.

As we shall see, at least when combined with other considerations, structuralism is interested in formal descriptions of final states, even if in order to understand how these emerge, non-structuralist considerations must become paramount. In evolution for instance, the structuralist would be interested in the formal properties of flying that are common across all species that fly (because all are constrained by aerodynamics); whereas to understand the origins of flight, and the different evolutionary paths that converged on it, then other contrasting issues must be studied: evolutionary history, the fitness of each evolutionary step in its own niche, and so on.

Structuralism could have large implications for teaching: if you believe knowledge comes in lumps then presumably this tells you about what must be taught together. Furthermore it leads to views like those of Rumelhart & Norman (1978) that not only knowledge but the process of learning is lumpy. Whereas a non-structuralist must expect that all items are equally hard or easy to learn, a structuralist is likely to predict large differences depending on whether the learning requires a whole new structure ("restructuring"), modifying an existing structure ("tuning"), or simply slotting a new item into an existing structure ("accretion").

What Spiro defines as ill-structured domains, where no one structure uniformly applies, must be seen as an important third case: neither the pure structuralist case where the structure captures everything nor the opposite amorphous case. Many domains of human expertise, ultimately perhaps all, are partially understood in terms of several alternative structures all of which capture important partial aspects. An example would be the wave/particle duality before their synthesis by wave mechanics. Knowledge in such domains is organised primarily around cases, secondarily around the parallel structures.

However, whether knowledge is seen as structured or amorphous seems to be independent of the other central ideas about teaching and learning. For instance if you are a constructivist you could be either a structuralist or not.

2. Constructivism

[Key reading: von Glasersfeld 1984]

Constructivism is a term that crops up in diverse fields, but the root meaning discussed here is the sense used by Piaget and von Glasersfeld (1984). As we shall note below, in the end it has little to do with most everyday meanings of "construction". "Constructivism" labels a position in epistemology (the theory of the grounds of knowledge) which holds that there is never enough information for a person to be sure of the beliefs they hold. This is basically because, although experience can falsify our beliefs, it can never prove them true. It thus belongs philosophically to the tradition of sceptical arguments such as those of Hume, Popper, and Kripke (1982). A consequence of this is that a person's beliefs are not a simple reflection of an external reality, but result from the interplay of two factors: constraints from external events (which may disconfirm beliefs), and the individual's prior mental contents i.e. the history of their mind.

The proposition that learning depends in part on the learner's prior mental contents is the key one for education, not least because it relates directly to the findings of work on prior conceptions: their prevalence, variability, and resistance to instruction. Thus whatever its philosophical merits, constructivism seems relevant to the TLP in higher education. Note however that it is, and needs to be, a dual theory that covers the range between two extremes: cases where transmission works and the learner's beliefs converge closely to external reality or teachers' concepts, and cases where little convergence occurs and the contribution of prior mental contents dominates. Teachers are familiar with cases of both kinds. These two cases can seem incompatible, and we may feel that constructivism is like the wave/particle duality in quantum mechanics, but really it is a theory of the interplay of two factors, not of oscillating between rival explanations to suit each case. In this way it is closely analogous to the theory of evolution, where the two factors corresponding to external inputs and prior mental contents are adaptation to the environment and the organism's evolutionary inheritance. Because constructivism is a theory of the interplay of two factors, it does not have a strict opposite, but rather opposes both of the rival theories claiming that one factor overrides. Objectivism or realism claims that external inputs determine a person's beliefs, and their opposite is extreme relativism: that only prior mental contents matter, there is no truth or knowledge, and any belief is as valid as any other. Constructivism is the view that both are partly right about the factors affecting beliefs, but both are equally wrong in going for a one-factor theory. In the constructivist view we can never be sure that we have arrived at the truth, but beliefs are not all equally valid: they have to fit external inputs (so far), and we can expect convergence as we encounter more evidence.

Another relevant comparison is with communication. Technology from Morse code to ethernet and the error correcting codes used to transmit data from spacecraft to earth is full of examples of communication using fixed codes built in to both parties in advance by human engineers. This is the situation assumed in Shannon & Weaver's information theory, and also in most of linguistics in notions of adult competence. Although human communication certainly includes this case of a fixed, pre-arranged language it also includes cases where the code is set up or changed during the interaction. This must be the case of course for each person learning their first language, before there is another shared language in terms of which a new one can be defined or explained (Quine 1969). This case is much less understood, both theoretically and intuitively, but it must involve participants attaching new symbols to existing mental contents; and inherent uncertainties in whether the same links have been made by both parties. A general theory of communication must cover both cases and all mixtures of the two, just as constructivism covers all interactions of external inputs and prior mental contents. This is directly relevant to teaching, particularly of new concepts where by definition the learner does not in advance possess either the words or the meanings. Substantial difficulties have been found to exist in practice for the vocabulary (learners often do not understand the same meaning as the teachers for terms such as "volatile" or "work" — Cassels & Johnstone (1985)), for finding shared experiences to which to link the concepts, and for the concepts (where prior conceptions often resist replacement by the new "correct" concept).

The best summary of constructivism, then, is the view that learning, in fact thoughts in general, depend on both current "input" from the situation and on prior mental contents. This in fact applies not only to learning new concepts, but to communication, to explanations, to tasks, and to motivation. Communication and interpreting what people mean by what they say cannot always depend on previously fixed codes. What will count as a good explanation for you depends on what you do and don't already know that can act as an anchor for the new explanation. Deciding what task is being demanded is equally a matter of interpretation based on experience. The work on deep and shallow learning shows how that matters in learning tasks. It appears even in apparently very simple tasks: if you are training someone on elementary word processing and give them a specimen letter whose format they are to reproduce, people vary in what aspects they take to matter: from printing blemishes, through the question of setting the margins to the nearest two inches or the nearest millimetre, to not noticing whether the text is right justified. Similarly with motivation, where external incentives (the contribution of the current situation) are important, but so are "intrinsic interests" i.e. the way particular learners have a prior interest in particular subjects.

Constructivism warns us that teaching cannot be taken as the sole, or even main, cause of learning; although equally teaching, in the role of "external constraints" or input, is important as a co-cause. Thus neither teaching as transmission nor learning as autonomous discovery are general theories compatible with constructivism, although both are possible in special cases. That is why I have been referring to the TLP — the teaching and learning process — as the basic complex we are concerned with; and is fully compatible with my opening assertion (also supported by painful experience) that teaching cannot cause learning in any direct and reliable way. Thus not only can you not reliably bully someone into learning because however frightened they can only interpret what you assert in terms of what they already understand, but equally you cannot share your knowledge: it will always be differently constructed (and occasionally better). Violence, patronising wisdom, kindness, and empathy are all equally unreliable in securing accurate transmission: at bottom constructivism is a bleak philosophy that denies the possibility of knowingly sharing knowledge.

It is hard to find a good discussion of constructivism for (higher) education. (The appendix to this paper offers a longer discussion.) The philosophical position is well discussed in the chapter by von Glasersfeld (1984), and more widely in that book (Watzlawick; 1984). The essential philosophical strands of Piaget's thinking are laid out in Piaget (1970/1983), although constructivism is woven in with structuralism, activity orientedness, and Piaget's development of these ideas in concepts like accommodation. Authors in the education field are mostly even less clear about the relationship of constructivism to other concepts they believe in; and perhaps because of that, many other authors do not mention constructivism even though they clearly accept the key notions. The most common misuses of the term are to mean a) Piaget's views, b) whatever the author believes about learning ("constructivism" is used to imply "good"), c) constructivism as a simple opposite of a transmission view of teaching (cf. Jonassen; 1991), which is actually a mistake as implied above, besides being contradicted by our everyday experience of cases where transmission works.

Finally we should note that "constructivism" is not turning out to be a particularly felicitous label for the view of learning it denotes. On the one hand even non-constructivists will agree that learners construct: even a tape recorder can reasonably be said to construct the magnetic patterns on the tape based on sound waves in the air. In other words any kind of information processing view of the mind that allows for mental representations implies some metaphor of construction, but not the philosophical view of "constructivism" which is meant to denote only mental construction to which the learner has made a contribution that depends on their individual mental history. On the other hand, constructivism does not imply, and has no connection with, some of the strongest intuitions invoked by the metaphor of construction: that reflection and effort after understanding are important, and that education is or should be about restructuring, reflection, deep learning. This credo may be true in my value system and yours; but it is not constructivism. These ideas relate to other, independent, concepts: restructuring is closest to structuralist notions of mental schemas and of accommodation, reflection is one of Laurillard's TLP activities, deep learning is addressed in Marton's work. Thus the metaphor of construction fails to pick out the defining feature of constructivism in an educational context — the role of the learner's prior mental contents as one major determinant of learning — while encouraging us to slide into thinking about other, independent, aspects of learning.

3. Activity orientedness: being economical with your learning

Does doing have an important causal influence on learning, or only vice versa? Must the learner be trying to do something with the material to be learned, or not? Some constructivists (e.g. von Glasersfeld 1984) derive activity orientedness from constructivism, claiming that constructivism entails activity orientedness; that what is learned depends on the learner's activity i.e., on what they do and on what they want to do. The argument for this, inspired by Piaget, is as follows. To draw conclusions from new experiences means categorising things and events, and such categorisations start from what makes a difference to things you care about. If all you care about is throwing, then a mug and a ball are much the same though a heavy mug and a light mug are importantly different; whereas if you care about drinking then mugs and balls are totally unrelated.

I don't think it follows. The argument only works if you add the additional tacit premise of learning being economised on. Otherwise learners would make all sorts of other generalisations from their experience, based for instance on the perceptual properties such as colour available to them. In the case of people being economical with their learning, the only prior mental contents of a learner would be strictly business: to do with their goals. There is evidence both for and against this. Against it is all notions of curiosity. It is not obvious why it is in any agent's interest to fail to remember: there is no reason to think that learning takes up energy or any other resource that is limited: memory does not seem to have a finite capacity. On the other hand, we have found cases where people do fail to learn aspects of artifacts such as user interfaces where perception reliably delivers the information whenever it is needed.

However there may be a more basic reason for economy, and so for selective and limited learning, than a kind of laziness: the combinatorial explosion of possible things to learn from any event. Suppose I hand you a pebble from the beach and say "learn this stone": What will you learn? Is it laziness that will make you learn different things from other people, and prevent you from learning it all? A pebble is a small, simple, very ordinary object: how could it be a problem to learn? All learning in fact involves generalising in some way: first selecting out some property then treating that property as going beyond the event itself. Even a tape recording, the crudest kind of "learning", does this: it loses the time and place and the loudness (which depends on distance from the microphone, volume settings etc.) of the original sound, but it preserves the main features we need to recognise speech and music. When people learn they similarly select and abstract, and they must do so because there are an infinite number of possible attributes and relationships presented by each object and event. People are flexible and can choose what to notice and perhaps learn; if they are energetic they may even learn a lot: but it will only ever be a tiny fraction of what was there. If we consider how that selection might be made, we see it could be described as being made in relation to tasks: to what use the information would be. If you are asked to throw the stone, to use it as a token in a game of Go, to classify it geologically, to describe it to someone as a decorative object for a mantelpiece or collection, these would all immediately direct you to different properties. You would not have to actually perform these tasks: given that you had some understanding of the tasks, just considering them would be enough to direct your attention and learning. That is why I refer to this selectivity as "activity orientedness". Another feature that is obvious from this example is that if the situation does not explicitly demand a task, you are likely to invent one by drawing on your experience, and that people are likely to show differences because of this.

Even learning that appears quite passive is in fact oriented to some task or activity in the above sense. For instance Pavlov's dogs learning to salivate when a bell sounds are oriented to eating. Chess experts memorising board positions at a glance do so in terms of structures useful (i.e. oriented) to playing the game, and cannot remember unlikely and impossible positions well. When we remember a person, we remember facial features that will allow us to recognise that person and distinguish them from others: we do not remember the lighting, the angle at which we saw them, we are not very likely to remember their clothes unless they are distinctive, and it is even quite common with people you know really well to fail to notice the removal of a beard or moustache or several inches of hair. This all applies to academic learning, where usually students are not asked to and usually do not learn page numbers and font sizes; they may or may not be required to learn particular words and phrases; and almost certainly will be intended to learn the "meaning".

It is also important to recognise that learning for some tasks can incidentally equip you for others; this is a highly non-reciprocal arrangement as the work on deep and shallow learning explores.

In summary, the argument put forward here is that it is wholly impossible to learn all aspects of even the simplest object or event. Consequently learners must and do exert choice in order to impose a drastic limitation on what they learn. As an extension of that, they may decide to learn nothing, either to save effort or because there is nothing new to learn (relative to any of their goals) e.g. "just another pebble on a beach with millions". We can describe this choice in terms of being oriented to some activity that would use what is learned. Most frequently people learn nothing (dismiss the item as irrelevant to any of their concerns); next most frequently they will orient to some single task; however sometimes they will orient to several tasks and learn several unrelated aspects (e.g. remember the lecturer's shirt in connection with making jokes later, as well as remember something of the lecture topic in connection with passing an exam). There is nothing impossible in learning multiple aspects, related to multiple tasks, although especially in adults economy of effort in learning as in all activities seems to make this uncommon. However a person or child (or young bird) may in fact digest experiences by trying out a small set of alternative (even standard) tasks: eating, rejecting, attacking, etc.

Although this selectivity is open to choice, how learners choose will in general be constrained by the two kinds of influence described by constructivism: inputs from the current situation (e.g. exam pressures, teacher requests, a need to get this equipment to work), and the learner's prior mental contents (intrinsic interest in some things, experience of the kind of thing examiners generally want you to recall). While constructivism says that all kinds of prior mental contents are important, activity orientedness says that goals / tasks are uniquely important to learning.

Implications for the TLP

The distinct view labelled "activity orientedness" is that goals held by learners strongly influence learning, where "goal" is to be interpreted as striving to achieve some task such as answer a question, solve a problem. An implication for teaching would seem to be that it is desirable to communicate early the tasks with which material will be tested, the tasks for which it is useful. This tells learners what to take away from the material. It should avoid the waste of time when learners sit listening to a story without yet seeing what the point is, what

to attend to. However the work on deep and shallow learning seems to contradict this, and certainly shows that there are pitfalls. In that work, all attempts to suggest specific tasks, even that of reproducing "deep" explanations, in fact increased the prevalence of shallow learning. On the other hand that body of work is probably the most direct evidence that what is learned is strongly affected by the goals the students relate to the material, and that very different learning outcomes from the same materials and apparently similar students can be explained in this way as a difference in the goal grasped or adopted by the learner i.e. a difference in the activity anticipated or associated with the learning.

Piaget

[Key reading: Piaget (1970/1983)]

As much of Piaget's work concentrated on pre-school children, it might not have much relevance to a theory of the TLP in higher education: quite different issues might be the dominating factors. In fact the limited relevance of going from the study of young children to university students surfaces in the topic of motivation. Piaget had little to say about this, but its importance is all too salient in higher education, as we discuss later. However the three factors picked out here as the most fundamental grounds of his views — structuralism, constructivism, and activity orientedness — are all of central importance to teaching in higher education. Indeed some of his theoretical developments from these foundations, such as his concepts of assimilation and accommodation, may be useful too. For those interested, what better starting point than Piaget's (1970/1983) own summary "Piaget's theory".

A point to be aware of, though, is that while Piaget has had the biggest impact on developmental psychology and education, his own primary interest was philosophy. He studied the origin of knowledge in children, not for what that showed about children but for what it showed about the nature of knowledge. Thus the interest and purposes that shaped his research and writings are exactly opposite to those discussed here: we are interested in what philosophical issues like constructivism may have to say about the TLP and hence about how to design better teaching, whereas Piaget was interested in how observations on how children learn (or at least develop) might constrain philosophical arguments.

It seems useful here to separate out the logical strands of structuralism, constructivism, and activity orientedness, because it would be entirely possible to hold any one or two without the others. For instance you could be a constructivist, but believe that learners construct all beliefs consistent with their experience (past and present) rather than only those related to their activity. These beliefs would be diverse and history dependent, but not limited by doing. Or you could believe in activity orientedness but not be a constructivist, and believe that learning was the reception of knowledge from a teacher, but only reliably retained in connection with specific activities. (Learning a second language fits this viewpoint rather well.)

However for someone who holds all three their relationship is something like this. Structuralism is concerned with the final state towards which an organism converges: the knowledge a child acquires in common with others, the adaptations to an environment that even quite different organisms may exhibit. Constructivism is about the process of convergence, about how organisms may be quite different initially, about the dependence on history and how this interacts with the pressures towards adaptation. Activity orientedness is the additional claim that activity is the main test and filter of learning, comparable in evolution to the claim that all that matters is how many descendents an organism leaves: these are the proximal causes of convergence, and may modulate or even distort the process compared to the predictions you would make solely on the basis of adaptation type arguments about what the process is "meant" to converge towards. Thus in evolution sexual selection and other "selfish gene" phenomena lead to traits that are not adaptations in the older functional sense; and in learning, the tasks learners limit themselves to may prevent them from learning the topic "properly" from an educationalist's viewpoint.

4. Forgetting, reconstruction, and teaching multiple links

A basic fact of human memory not accounted for in the above, or indeed in most learning theories, is that people forget. This is in contrast to machine memories, and most simple ways of thinking about memory as a container. It is the more puzzling in that there is no evidence at all that people ever run out of memory space and need to forget to make room for new material.

When someone has partially forgotten something, they often reconstruct it partially or fully from other material. In fact Bartlett (1932) showed that reconstruction was a fundamental and pervasive characteristic of human memory. Because of our tendency to forget, retention depends to a considerable extent on the number of ways we can reconstruct something: on the number of rules or links we learn to an item. We are thus led to a teaching principle, whose value has been shown repeatedly in many ways: material that is linked in multiple

ways to other things the learner knows has a better chance of being learned, and will be retained longer. If we did not forget, then multiple links would not be important except perhaps to allow people a choice of alternative links to learn: once any link was learned, then the item would be permanently known. In practice there is overwhelming evidence to the contrary, and learning multiple links is very important.

This teaching principle of multiple links and the learning principles of forgetfulness and reconstruction are independent of the other factors discussed. It is clear that forgetfulness is not related to structuralism, constructivism, or activity-orientedness. The value of learning multiple links is independent of structuralism: reconstruction can be by rules (the structuralist case) for instance if you forget an entry in your "times tables" then you can reconstruct it using the rule of repeated addition e.g. 4 times 7 can be calculated as $7+7+7+7$. However reconstruction is just as important for non-structuralist cases: semantic nets, or completely arbitrary associations e.g. the standard mnemonic techniques for retaining arbitrary items is to link them mentally to shapes, colours, positions etc. Constructivism is equally neutral: redundant links are just as useful whether they are to personal items others do not have, or to standard items shared by everyone. Indeed many supposedly constructivist teaching tactics in fact rely on making connections between the new material and something that all learners uniformly know already: their value is in making extra links beyond the general meaning of the words, not in attaching meaning where no comprehension existed. Multiple links is if anything against the implications of activity-orientedness insofar as that implies economy of interpretation of new material. Teaching that works to set up multiple links will be working against the tendency to economy; although conversely, given the desirability of multiple links, the need to push for this specifically rather than expecting it to happen automatically is predicted by activity-orientedness. One could thus see multiple links as a teaching intervention necessitated by the combination of the tendency to forgetfulness, our ability to reconstruct to counteract forgetting, but our tendency in learning to limit processing (inferences, discovering multiple consequences and views) in line with activity-orientedness.

5. Concepts vs. Tasks as the content of learning

Activity-orientedness does not imply a tendency to learn procedures rather than concepts or declarative, factual knowledge; at least not in adult learners. As Piaget argued, facts and concepts are generalisations that emerge about actions, and learning them allows one ultimately greater power of action over the objects. This effect permeates even perception and memory. For instance the demonstration that chess experts see the board in terms of chunks that allow them to understand and memorise boards vastly better than novices only applies to those board positions that actually occur in normal games and so are useful. Thus what looks like a perceptual skill, or an ability to learn static facts (board positions) actually comes from a prolonged and practiced orientation to playing chess.

These representations are factual in the sense that they are neutral between a number of immediate tasks (playing chess, memorising board positions) — that is why they are abstractions worth having. But like all representations, they are not neutral between all possible tasks but on the contrary are optimised for some and not others. This again is a very general fact about representation, familiar to all computer programmers.

Another thing that follows from this is that whenever more than one schema or representation can be applied to the current input, which schema is applied will determine what the person notices and hence what they could learn (remember) about it. For instance if I show you the number 91627 you will probably just think of it as some serial number, and find it hard (and uninteresting) to remember. If I suggest it is a sum of money, you might recall it tomorrow as "about £91k" or "\$91k". If I hint that it contains a hidden sequence, still more that it is to do with squares, you are likely to be able to remember it easily. (See Katona; 1940/1967 for more.) Different schemas in the mind of the learner can have profound effects. Each schema originated in different activities, which in the long run affect concepts and factual learning as much as action. In the short run, this is mediated by which activity is selected as the relevant orientation, and this affects not just what the learner tries to do, but how they describe any object. Descriptions, then, will depend on both current input (what the situation suggests is the relevant activity to orient to), and the learner's prior mental contents (the selection can only be between things the learner knows about: non-chess-players do not have the option for describing and memorising board positions in the way experts do).

6. Learning shared descriptions: the socially distributed nature of knowledge

The philosopher Hilary Putnam (1975) pointed out that what almost all of us mean by "gold" or "water" is "whatever an expert chemist means by them". Although we can rule out many things as not gold or not water by their appearance, in cases that cannot be easily resolved like that we refer our decision to others: much of what we mean and in fact of our knowledge is ultimately grounded in a social network, not in our own

independent grasp of the world. It is not that our personal and perceptual experience is unimportant, but when in doubt we tie our ideas to other anchors — and in particular through language to what other people know. Our knowledge, and therefore what we aim to teach others, is therefore mainly something that is socially distributed and not something which individuals have separate and personal copies of. A further consequence is that a crucial part, in fact the most crucial part, of what we learn and teach about most things is not its nature but what other people call it. This social distribution must necessarily depend upon a public shared language: words like "gold", "water", "uranium", and other notations, terminology, formalisms.

In the light of this observation we need to review an important trend in education, particularly science education. In the last few decades a lot of work has been driven in one way or another by a feeling that a lot of education turns out students who can pass exams on paper without ever relating what they have learned to their own experience: e.g. chemists who can write formulae but do not know the colour of copper sulphate. This is clearly undesirable: we do not want to say such students understand their subjects as well as those who do connect their academic knowledge to their own experience. However in the light of the socially distributed nature of knowledge we must recognise that knowing public names and descriptions is more important than personal experience, and if a student has only one then the exam knowledge of descriptions of concepts and substances is more important than personal experience, even though both would be preferable. After all, if you know the name of something you can use reference books, ask people, and quickly draw on society to discover more.

The main content of all academic subjects in all disciplines is in fact descriptions, not sensori-motor experience or private knowledge, and this is not an accident or due to bad educational systems, but is a direct reflection of what most knowledge is and how it is useful. Academic knowledge is primarily knowledge of descriptions. In higher education at least, activity orientedness thus has to be interpreted primarily in terms of activities with descriptions, not physical actions on the objects of study (e.g. molecules, 16th century France). There is also a secondary level of personal action and experience which needs to be learned and related in good education: if you are a chemist you should know what gold looks like, if you are studying literary criticism you should know what it feels like to be in the audience of a theatre, if you are a doctor you should know what "extreme pallor" looks like in a patient not just that it means unusually little blood near the skin surface. Laurillard (1993; e.g. p.23ff. and p.58ff.) makes this distinction, and makes it central to her model of the TLP. However because of the importance of the social distribution of knowledge, especially academic knowledge, we should not only recognise the central importance of the descriptions level of learning, but that it is not grounded in the level of personal experience. Instead, the dimension of personal experience is only an adjunct and likely to be quite variable. No-one can recognise momentum or gold by sight with complete accuracy. Experts will do better than students, students than those who haven't studied it at all. However each may tend to use different perceptual characteristics (colour, density when picked up, or surface scratches for gold), partly because none of them take perception to be the final grounds of identification.

Because the level of descriptions is the most important, it is not surprising that many learners do not naturally and automatically go beyond learning the words. The principle of economy implicit in activity-orientedness will make it natural to stop at the surface level of the names and descriptions, particularly since this captures so much of what it is to know something. There is no absolute standard of whether you understand something: Putnam's argument shows that very few people understand the meaning of "water" in any deep sense. Recognising a word and knowing how it is used constitutes all the understanding we have or need of many things. Constructivism does not contradict this: a learner's construction of a meaning can be just its connection to cases of hearing it used, and if these cases are not representative then there will be a "misconception" whereas if they are, then no misfit will occur. Understanding thus need not include connections to personal sensori-motor experience.

If teachers decide that learning connections with personal experience are important, then they will need to make this a distinct learning objective and promote activities using the connection. They may decide this is additionally worthwhile as a way of creating multiple links and so longer retention.

Teaching public, common descriptions is necessarily didactic: done by telling. It is obviously nonsensical to expect learners to re-invent the same words as are commonly used by others. This is not to say that any actions (e.g. the steps of arithmetic calculation) should be taught by telling rather than by the learner doing. This combination can lead to contorted dialogic practices (Edwards & Mercer; 1987).

7. Schooling

A common assumption is that knowledge is what can be tested for in an exam: recall without prompts of any kind, and that the kind of learning done for tests of this kind is the one and only universal human learning mechanism. Cross-cultural work however suggests that proficiency at this is a function of years of practice in school (Scribner & Cole; 1973). What are the distinctive characteristics and effects of informal, formal, and school-based learning? Scribner & Cole argue that school-based learning is decontextualised learning: learning out of context, general, abstract, rootless. It is about learning to attend to surface forms, the language itself, and to rules not effects. Margaret Donaldson points out that in normal conversation, people care only about what the speaker means, intends to convey; and not about how they say it or what the language could mean in other contexts. But school is about changing that orientation, and developing skills in literal meaning, and rule based interpretation. That is why it is unnatural. In fact, schooling is training people to do shallow learning (for recall tests, ignoring context and meaning) that they couldn't naturally do. Cole & Scribner (1977) argue that the recall tests that seem so natural to most psychologists are in fact testing only a skill specific to schooled cultures. In other cultures, people do badly at most recall tasks, and certainly at tasks about learning for recall; while they will do well at recognition tasks.

However such training, as divorced as possible from meaning and context, brings many benefits. It is in fact an extension of the importance of the descriptions level of knowledge. The point of such decontextualised learning is a) to permit language to cross time, space, and contexts better, and so to do cultural and information transmission; b) the great power of formalisms is exactly about doing rule based manipulations of the surface form without considering (using) the meaning, then re-interpreting the result, and getting results you wouldn't otherwise get. You may check a calculation by comparing the answer with all you know about its meaning in context, but the power of the method is that you manipulate symbols without any reference to the context during the manipulation. Thus you can develop rules for adding apples that are the same as for adding oranges, even though adding them together makes no sense.

Similarly, when you go to a doctor, you expect them to recall reliably not only the most common drug treatment for your problem, but also the side effects of it and the situational factors that would make that treatment undesirable. These facts are often unexplained theoretically, and certainly the theories of the operation of each drug need not be known by each doctor. Thus the crucial content of education for the most important professions has a large rote learning component, and the education and assessment methods used reflect this.

It follows that concern about "shallow" learning, where pupils pass exams yet fail to understand the connection between the concepts and vocabulary they have learned and the real world fails to acknowledge that the chief aim of schooling is to suspend disbelief, disconnect from the normal effort after meaning, and attend to the surface forms of representations. Shallow learning is not the product of immature pupils, but on the contrary, of well schooled pupils who are applying the basic educated technique. If we wish to argue that they are doing so in an unfortunate way, that reduces the learning benefits they might otherwise enjoy, we must construct our criticism in a framework that acknowledges the great benefits of such learning approaches in many case.

In summary, there are 3 points:

1. Schooling is about a special type of learning — willed learning for recall — that is not general learning. Schooled cultures develop special abilities at it; other do not.
2. The content of learning is mostly descriptions, not material actions. Again, this is not a general fact about human nature, but a consequence of the great degree of knowledge specialisation and hence of the social distribution of knowledge in our culture.
3. There is a special power to learning and manipulating surface ("shallow") aspects of descriptions, as opposed to their meanings. Schooling is centrally about developing an otherwise abnormal ability to do this and so to suspend our otherwise pervasive effort after meaning.

There are problems if this is all that is noticed about learning; but in remedying the bad effects of taking this as the whole story we must not forget these basic facts about learning in our culture.

8. Motivation

It is very obvious, particularly in university teaching, that learners are often highly selective in their efforts: what they learn often depends on their motivation e.g. being told that this topic will be examined. In fact whether or not students believe a topic will be examined is a much bigger factor for determining learning than any other, certainly than the quality of the teaching. This is not only a matter of common experience, but has also been found in experiments. It has vital implications: for instance testing a teaching method in any situation other than one where students will be assessed with the same consequences as in the final situation is unlikely to predict its effectiveness in the intended situation of use. Another implication is that the TLP is quite different with and without that external compulsion or motivation. Without it, learning is not zero but it is much less,

depends much more strongly on the quality of the presentation (Marsh; 1984, 1987), and what gets retained is often not the message that was intended but only some aspect of the presentation, as advertisers have found.

Many learning theories have no place for motivation, and thus implicitly claim that it has no effect e.g. behaviourism, connectionism, SOAR. For them, learning occurs uniformly, inevitably, and independently of the intention of the learner to learn. Similarly neither constructivism nor Piaget have anything to say about this issue of motivation, perhaps because it may not be important for infant learning, but it seems to become increasingly important with age, so that 40 year olds (say) learn little that is not directly relevant to their work, although they are often very efficient at picking up material that is and show much greater powers of concentration than students do. Thus it was often said that managers were either unable or unwilling to learn how to use computers, and what a "problem" this was (certainly to computer salespersons). However when tools appeared that actually directly helped managers with one of their tasks (financial planning), suddenly this "resistance" seemed to vanish. This is not because those tools are instantly learnable with no effort: rather, because managers like most adults will invest effort in learning if it will be clearly repaid, but not because someone else assures them that it will be good on some general grounds. Undergraduates on the whole seem to be more like executives than like infants in their approach to learning.

The key feature is that effort makes a big difference to whether something is learned and retained. Learners must necessarily limit their learning to one or a few aspects of what is presented to them, but adults show a strong tendency to further limit whether they try to learn at all. Consequently motivation, in the sense of willingness to make an effort, becomes a big factor in the TLP. Theories which omit this are so seriously incomplete that they cannot account for the major facts and activities in education e.g. exams, revision, dismissal of material that will not be examined. Many teachers have been influenced by such theories, and are therefore continually surprised and hurt students' manifest attitudes.

Motivation thus fits with activity orientedness, but is distinct from it. If learning depends on the tasks that learners focus on, then it is unsurprising that learning will depend in part on the willingness to perform tasks. Effort alone however is not necessarily enough as has been shown experimentally: perseverance in a task that does not in fact promote learning (e.g. endlessly re-reading a poor set of notes) does not help. Later sections discuss the different effects of different kinds of task. However it is clear that motivation will be one important factor in general. In other words, the general proposition of activity orientedness is that learning depends on the tasks to be served by the learning. One issue will be what task is anticipated or engaged in (whether it is a material task or one of description). Another will be whether and for how long this orientation is pursued by trying to learn or by practicing the task. "Motivation" is the label for a person's willingness to pursue it.

In other words motivation to learn is like motivation to acquire money: it is not sufficient, it is not even wholly necessary (cases occur of both money and learning being acquired without effort), but it makes an enormous difference nevertheless. For learning, in the absence of external motivation learning some things may still occur because of internal or intrinsic motivation. If I just say to you "learn this stone", then you are not very likely to learn anything unless either I am in a position of authority over you, or you happen to be a geologist. In other words, for any topic there will be a small minority who find it interesting anyway, but the great majority will not learn without an external motive.

Thus we may distinguish two classes of motivation: external organisational pressures, and internal motivation or "affective reaction" to a teaching and learning situation. It is clear that a learner may find some particular material intrinsically interesting and that this will make it much more likely for them to learn it. This probably explains not only the existence of some learning in the absence of external motivation, but also why it varies so widely between individuals depending on the content. In this it seems essentially the same as motivation for activities other than learning, where under the name of amusement, hobbies or sport one minority will do things such as sitting in rooms polluted by smoke and noise that others would regard only as a health hazard, or cycling up mountains that others would regard as cruel and unusual punishment that even criminals should not be subjected to. This contrast between external and internal motivation is essentially the same as in constructivism: the contrast between what the situation gives or "demands" and what the individual brings as part of their prior mental contents, both having a major effect in at least some situations.

Motivation in the sense of willingness to pursue activities is thus an important factor in the TLP, with two major subclasses: external organisational motivation (e.g. qualifications dependent on passing exams), and internal (e.g. intrinsic "interest"). It is clearly consistent with structuralism, constructivism, and activity orientedness. However it is logically independent of them, as you could believe in those factors while believing that motivation was unimportant as indeed it may be for infants, or conversely that motivation mattered but not the other issues. It seems to be of great importance in higher education, and probably for all learning after the age at which a child recognises that conscious efforts to learn make a difference.

Deep & shallow learning

[Key reading: Marton et al. 1984]

The characteristic experiment on deep and shallow learning goes as follows. An exposition is given to each student e.g. a chapter of a textbook, which they study knowing that they will be asked about its content. Open-ended questions are then asked, and the answers recorded. The researcher reads these over repeatedly, and comes up with about 4 types of answer. He then defines a coding scheme and checks that other people can categorise all the answers according to his criteria. (Note that this can be said to be one of the few truly constructivist measuring schemes, as the categories are derived from the variability actually found between learners, not from teachers' views of the material.) These categories usually, though not always, can be put into a single order where answering to satisfy a category implies satisfying all categories "above" in the ordering. The coding scheme can then be referred to as deep (the best) through to shallow. The essential phenomenon is that the same instruction gives rise to very different learning outcomes, of very different pedagogical value, for different students. Deep learning is associated with long retention of overall meaning, shallow learning with brief retention spans.

Learners are interviewed about how they were trying to learn, characteristic differences in what they reported emerge, and these too are categorised. These correlate with the learning outcome categories. Looking for facts and trying to memorise the text leads to shallow learning, while trying to understand the point of the article leads to deep learning.

Manipulations can be done to try to affect the kind of learning that occurs e.g. by in-text questions. By asking shallow questions, shallow learning can be induced in far more students. Thus the type of learning is not a fixed characteristic of individual learners (such as their prior mental contents, or metacognitive ability). As Svensson (1984; p.66) says, the quality of learning is the quality of the interaction between learner and material. It is a relational phenomenon depending on both learner, material, plus other "situational" variables since a learner may treat the same material in different ways. The differences in outcome can be explained as due to differences in what the students were trying to do. "What we found was that the students who did not get 'the point' failed to do so simply because they were not looking for it." (p.39, Marton et al. 1984). The issue emerging here as controlling the quality of the learning is the interpretation by the learners of what their learning task is. When you tell people to learn something e.g. a stone, a chapter in a textbook, this is not an unambiguous statement: it is interpreted by different people in different ways even within an apparently uniform academic culture.

9. Selection of immediate task

We can conclude from this that the type of learning depends on the task orientation adopted by the learner, and so the schema they bring to bear on the material. The type of learning depends strongly on the type of task the learner can imagine doing with the material, or knows or believes is required. The form of learning task selected by the learner depends on several factors. 1) On what they believe they will have to do with the material — this is affected by practice tests, in-text questions, experience of apparently similar situations. 2) On learning skills and schemas that the learner already possesses: if the learner doesn't know how to understand the text then obviously they can't. 3) On the material: on whether it is amenable to their schema or skill. Thus bad material can reduce students to shallow learning. So can making it clear to them through practice that they will only be tested on shallow aspects. Note that this phenomenon is constructivist (though it has extra elements as well), applying here to the immediate task rather than to conceptual material: that the outcome depends on both current situational input and the learner's prior mental contents.

The importance of task selection by the learner is worth treating as a separate issue or principle, although it has strong links to previous principles. It is a case of activity-orientedness: of learning depending on the anticipated task. However it involves not only the anticipated task, but in addition the learner's method of learning intentionally in preparation for that task. It is also a case of the effect of learning multiple links, since deep learning causes better retention of shallow items than trying to learn them alone. However this effect is strongly non-reciprocal, leading to asymmetric effects that are not predicted by those earlier principles, but which are of great practical importance in education.

10. Learning as a task

From quite an early age, children become aware that they can make a difference to whether they can later remember something. By the time they are in higher education, this is a dominant feature. Learners want and try to learn, and this determines what they do e.g. re-reading texts, staring at particular items trying to learn them. Their efforts are not always very effective, but they are clearly an important factor in the TLP. The effects of motivation and of the anticipated task (task-orientedness) will be mediated through the learner's idea about how to learn that material.

11. Learning through doing

Although clearly some learning occurs without any actions being performed, it is also clear that important learning occurs through attempting tasks. One important aspect of this is when the original material e.g. instructions on how to find a person's office are clearly incomplete. Attempting the task then involves learning a lot from the environment, guided by the hints in the instruction. Later repetitions of the task will probably use more information that was learned from the first execution than from the original instructions. Another aspect is that performing a task often requires rearranging knowledge into a form that supports the task. This may not strictly be learning new material, but it will certainly allow faster performance in future, and also longer retention because of learning new links.

12. Feedback

It is worth noting that when a learner performs a task or activity, they may or may not they get feedback on the success of their actions: this constitutes an independent factor in the TLP. Benefits often come simply from task performance without external feedback. For instance running 10 miles will improve your stamina and probably make you faster the next time even if you do not find out how long it took you. Explaining a concept to a fellow pupil has been shown to produce conceptual advances even though no feedback from a teacher about its correctness is given, and the fellow pupil is less advanced. If you are trying to learn a poem by heart, you can usually tell whether your recall is accurate without referring to the book again. Most people can improve a piece of their own writing by revising it themselves, without external feedback. In other cases, feedback seems crucial. For instance if you are trying to learn to measure something, it seems essential to have an independent way of knowing how accurate your own attempts are.

In some cases, being able to do without external feedback depends on learners having internalised a method of judging their own products even though they have not mastered a method of generating them correctly every time. Thus for instance, people can often judge their own writing even though they still need to revise what they write; they can tell whether they have produced a valid mathematical proof, even if they cannot generate one whenever they like. However when a domain is entirely new to them, they will need external feedback as they learn to internalise these standards. More generally, constructivism warns us that communicating new ideas is not unproblematic: feedback provides more opportunities for improving the convergence between learners' and teachers' views. Moreover, constructivism implies that because of the individual character of learner's constructions, teachers cannot be expected to predict the misunderstandings that arise. Convergence cannot be teacher-driven, but must be expected to depend on learners generating material and having teachers give feedback on it.

Thus while learning does not by any means always depend on receiving external feedback, and simply performing a task does not necessarily entail discovering how successful you have been in your attempt, it is often very important. Receiving feedback is thus an important possible feature of the TLP, and a feature distinct from simply doing a task or being task-oriented. It might be thought of as a teaching-side feature that addresses the two learning-side issues of activity orientedness and constructivism; but is not essential in every situation because either the situation or the learners can sometimes supply it. The range of cases includes:

- a) No feedback of any kind; no learning. E.g. throwing stones at a target you cannot see.
- b) Internalised feedback: the learner can judge success and failure e.g. throwing stones at a target you can see, revising your own writing, practising playing a tune.
- c) Simple judgemental feedback from others: telling a joke and seeing whether they laugh, doing a maths calculation and comparing it to the answer in the back of the book; getting a grade or mark for a piece of work.
- d) Diagnostic, descriptive feedback from someone, that tells you not just (i) how well you did, but also (ii) what to change, and perhaps (iii) "why" i.e. introduce new concepts and explanations in relation to your efforts.

The more the learner already understands, the less they are likely to need more complex kinds of feedback; even though they may still have a lot to learn, they have internalised enough to be largely autonomous in their learning. However the diversity and unpredictability of learner constructions and reactions to teaching which constructivism warns us to expect implies that the most complex type of diagnostic feedback is likely to be a necessary component of teaching, no matter how well designed it is.

Note that it is not easy to provide feedback by computer; or rather, it is only easy to provide type (c) feedback for rather simple fixed-outcome tasks. However we must expect type (d) feedback to be important some of the time. Furthermore learning often depends on the type of task the learner engages in, so the question is whether fixed-outcome tasks are adequate for the kind of learning we wish to support. Since on the contrary deep learning is associated with tasks of open-ended explanation, even type (c) feedback (cf. grading essays) is a serious difficulty for computers, just as it is a serious load on the time of human teachers.

Feedback has been treated as a separate factor here for several reasons. It is more a teaching tactic than a learning feature. It is a not necessary condition for learning: when people read a newspaper or attend an academic conference they are likely to learn something, but they get no feedback at all on that learning which is pure reception. (Although conversely, this may be exactly why so many conference goers say they get far more from the personal conversations i.e. with feedback than from the talks.) If you believe constructivism then you can predict that feedback will often be necessary. However many of the most fervent advocates of the importance of feedback in the TLP have derived it differently e.g. from cybernetics, or from stimulus-response behaviourism (which is about as radically non-constructivist as you can get). We can also note that one view of feedback is as a separate source of new information for the learner, rather than a method for correcting faulty transmission and reception: instruction gives them an idea for an activity, but feedback gives information from the world about how to succeed. This fits with the cybernetic view, but also may be necessary for teaching at the experiential level where learners have to relate things to their own private sensori-motor experience which could not be done by transmission. However these cases do not fit, as constructivism does, with the use of type (d) feedback which in effect is new exposition at the level of descriptions, but tailored to the individual as in a personal tutorial. The notion of feedback is broad, and it is often worth considering the different cases separately.

13. Teaching by requiring activity

It is often claimed that learning depends on being active, on using what is learned in action. The issue is a part of the relationship of learning and doing; whether doing has an important causal influence on learning. Some people claim you can only learn through doing. There are maxims (quoted by Laurillard) such as "I do and I understand". On the other hand most academics believe they learn from books, not from doing things, and that that is the point of a university and has been at least since the time of Aristotle. Furthermore astronomers and historians, to name but two disciplines, exclusively study things they could never act upon (remote stars, past civilisations). People effortlessly learn facts about things they are fanatically interested in, or even just surprised by (e.g. a seven foot man seen getting into a mini). People sometimes remember things from a few striking TV advertisements. People learn from scalding their fingers (say), including from having other people spill stuff on them. What can we say about the relationship of learning to activity?

It is clear that activity is not in fact necessary for learning, so what difference, if any, does it make whether a learner acts i.e. performs the activity in question. As argued above, in HE this is most likely to be at the level of descriptions not experience; thus someone training to be a doctor, the question is whether writing an essay or working through an exercise is necessary for learning (the descriptive aspects of) a topic.

As will be clear from previous sections, the position I take is that activity orientation is important in all learning, but that a lot of learning does not require actual task execution as the examples above show. Nevertheless actual doing often improves learning for several reasons:

1. Actually doing a task gives us feedback on what we need to learn. Learning without task execution only works well if we are well practiced with the task in general, and probably have well adapted schemas into which new information fits. To some extent we can anticipate what we need to learn before doing it, at least for familiar kinds of task e.g. learning a route from being shown it once. But the more novel the material, the more it is essential to try a task to get internal feedback on what we should have learned and so on what we should learn now. Even if we don't get external feedback, going through the task is often informative e.g. trying out a story or a talk or argument in your head, doing lookahead in chess. Task execution is useful to learning for a similar reason to the importance of activity orientedness: processing limitations. Learning is determined by activity orientedness because we can't learn every aspect of an object or event: so we filter it through relevance to some activity. However we don't always do this correctly: we don't always notice what we need to resolve or discover from the material, but failing at a task, even just failing at trying to plan it, tells us what we need to know.
2. Getting external feedback (which of course entails doing the task) is often important, either because we may not have internalised success criteria adequately (what will make them laugh? is this the explanation they meant?), or because we may be missing part of the solution and unable to generate it without help.
3. Being sure we overcome laziness. People are often economical with their learning: they could learn without doing, but don't bother. Adults are like this in all areas not just learning; and young children are different in all areas: much more likely to run in circles, to think for no reason, to explore out of play or curiosity. So doing a task is a way of making sure that learning happens by testing for it behaviourally.

In summary, task orientedness provides the necessary selectivity to guide learning, and at least in adults much

of our learning is done by anticipating what we need to learn without actually doing the task. Indeed, the concepts of training and education depend on this. Furthermore learning for one task can incidentally equip you for others; this is a highly non-reciprocal arrangement as the work on deep and shallow learning explores. Again, training and education depend on this: on teaching component skills, and training on tasks that are similar to the target task (simulations, practicing emergency drills, etc.). Nevertheless, actually doing real tasks and getting feedback is the most centering thing: tells us most directly what we need to know, and whether we have yet learned it.

Laurillard

[Key reading: Laurillard (1993) esp. p.103]

Laurillard's model is a reconfiguration of some of the principles seen before. In the context of this workshop, it shows that much the same collection of recommendations about teaching can be repackaged in another way. As we shall see, she derives her whole model from: the importance of feedback, teaching by requiring activity, the distinction between the levels of description and of experience. She never mentions constructivism, but her approach is closely consistent with it. However she does not pretend, as some soi-disant constructivists do, that exposition has no place; and she is very clear about the importance to academic level of the level of descriptions and the tasks that go with them (writing, explaining, etc.).

Laurillard (1993, p.103) offers a "conversational" framework or model of the teaching and learning process. Her starting assertion is that particular media and particular didactic techniques do not generalise very usefully. However there are some generalisations about the teaching and learning process that (she argues) hold for all university subjects and teaching approaches. These are the importance of a universal set of categories of actions and interactions of the teacher and learner, all of which must occur for learning to occur satisfactorily. She lists 12 of these, which are diagrammed in fig.1 in slightly rearranged form. Our description of them here is a modification and expansion of her own very brief presentation in Laurillard (1993). The 12 types of actions can be viewed as being generated by two underlying ideas.

The first important idea is that learning depends on learner activities, and that (1) listening to a teacher's exposition (or equivalently, reading a primary exposition in a textbook) is only one of a set, all of which are essential for learning to occur. The second activity (2) is the learner describing the conception as they have grasped it so far. In a tutorial, the learner might do this verbally; an essay exercise might require it to be done in written form; a student with advanced learning skills might do it themselves when they write up their notes if, but only if, they compose those notes themselves rather than writing to dictation. The third activity (3) is the teacher redescribing the conception in the light of (2) i.e. what the student said or did; for instance, giving the student feedback, and re-expressing the material in the light of what did not seem to get across. A lecturer spending considerable time answering questions may be covering this, but if there are no questions, or if teaching is by a computer program that cannot understand student input and respond to it by alternative material, then this activity is not supported. The fourth activity (4) is the student redescribing (correcting) their original attempts.

It is obvious that stages 3 and 4 in general may repeat in an iterative loop. However treating them as distinct from 1 and 2 captures the essential element of adaptive feedback as opposed to merely trying again. Laurillard's metaphor, which she takes from Gordon Pask (e.g. 1975, 1976), is one of conversation between a single teacher and a single pupil. In practice the process is usually more distributed: not only multiple students per teacher, but a university student often interacts with several "teachers" on a given topic — e.g. a textbook author, a lecturer, a different person giving tutorials — and in addition peer interaction is often important e.g. a student can learn from the effort of trying to explain the material to another student. The conversation metaphor however is useful in emphasising the importance of the interactive nature of the process, and claims that learning is jeopardised when this is lost (e.g. only activity 1, or only 1 and 2, are supported).

The second major idea in Laurillard's model is that all of this occurs at two levels: the level of descriptions (being able to understand, manipulate, reproduce the public academic descriptions, including specialist terminology and notations), and the level of personal experience and action on a world. In science this will be both lab. work and the parts of daily life to which the theory applies; in literature it will be the experience (including the emotions) of reading literature or seeing plays performed (as opposed to the academic language of criticism). The model lists four corresponding activities (5-8) at this level: the teacher setting a goal (a lab. task, a problem defined in real world terms), the student acting to try and achieve it, the teacher or the domain itself giving feedback on the action, the student modifying their actions in the light of it. Again the claim is that if these additional four activities are not supported and performed, then learning will not really occur.

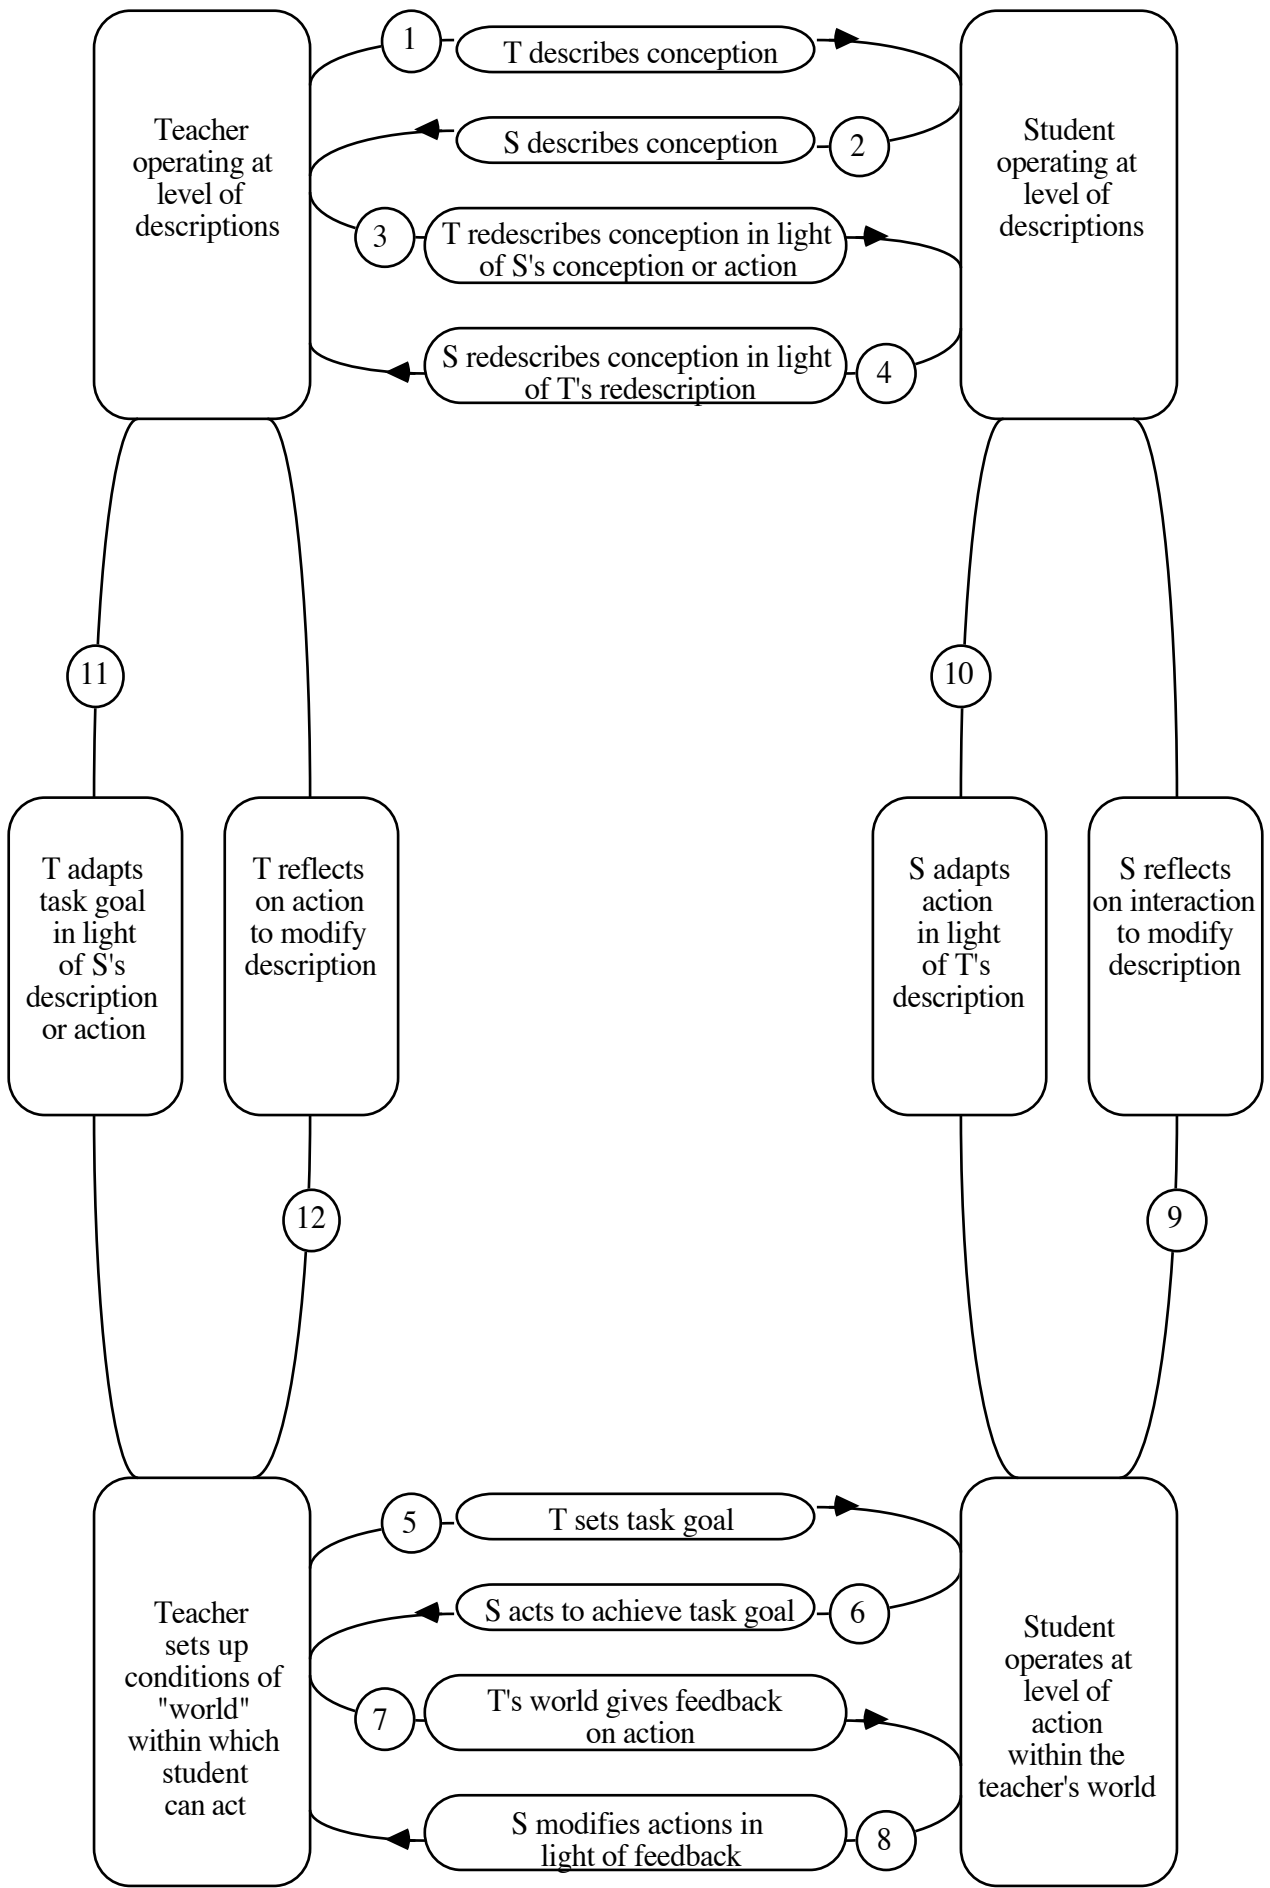


Fig. 1 . redrawn from Laurillard p.103 (fig.II.1) but with the numbers of the stages shuffled a bit.

These two levels then require a further four activities that link them: (9) the student reflecting on an interaction at the level of the world in order to modify their conceptual descriptions, or (10) conversely modifying their actions in the light of reasoning at the level of descriptions; (11) the teacher modifying the world task set to address some need revealed by the student's descriptions e.g. set a lab demonstration relevant to a conceptual problem the student exhibits; and (12) the teacher examining the student's actions and expressing help at the level of descriptions i.e. explaining what went wrong rather than showing what to do to correct it.

Teaching techniques

Above some teaching principles were covered: learning multiple links, preferring "deep" to "shallow" links i.e. encouraging trying to understand rather than trying to learn specific items, providing feedback, requiring learner activity. These are old principles, that although consistent with learning principles such as constructivism, cannot be derived directly from them, and have often been advocated independently of them.

Construction: manoeuvring learners into re-inventing the thing you are trying to teach. Of course this is not possible without enormous help by the teacher in shaping the activity to this one end. This tactic is used by behaviourists to shape pigeon behaviours, as well as by Papertians to shape child cognitions. Its benefit is not only teaching those you do not share a language with, but one more extra link. Possibly it could be argued to be a "deep" link that enhances understanding and therefore does more than most links towards retention. However this must depend on how constrained the construction process was: if very constrained, then probably the link has weak power as in that context the learner couldn't see any alternative; whereas when they need the link for reconstruction, confounding alternatives are likely to abound.

Puzzlement / conflict: The claim that learning depends on pre-arousing and linking to some question, some sense of the need for change, for an answer. This is a Piagetian idea, and appeals to our intuitions. It also seems to appeal to most journalists and documentary makers, who are forever beginning with some question supposed to tantalise you. However as usual it is easy to think of learning that occurs without it. Many of the most memorable facts cause puzzlement, not resolve it ("did you know, that the overall crime rate including terrorist offences is lower in Northern Ireland than any other region of the UK?"). Again, it is a particular form of the idea that more links are good for retention, the question being whether it is any better than any other link you could offer the learner e.g. this idea was believed by your hero. What it has in its favour in theory (I don't know of relevant comparisons with other links) is that a conflict is likely to be about understanding and so may prove a "deeper" type of link better at supporting reconstruction.

Reflection: Similar remarks apply: it appeals to our intuitions, it can benefit learning, but is it any better than other link-promoting tactics? It should be better, as "reflection" should stimulate not one but multiple link formation, all at the level of understanding.

Small group work: Piaget's idea about this was that it promoted cognitive conflict, but that has been shown not to be altogether a correct account. Another view of it would be that a learner is likely to switch often between generating and criticising accounts, rather like Laurillard's (Pask's) conversational model, whereas in most learning situations a learner is stuck in one activity for long periods. Multiple activities will promote multiple links.

Authentic tasks: This presupposes teaching by requiring activity, and asks what kind of activity. Naturally it had better be an activity you actually want to teach, not one generated only by the teaching situation. It is only interesting if there is a conflict between the activity required by the educational system (e.g. passing a written exam), and the "authentic" activity from other areas of life. The question then is whether training in the wrong task is nevertheless more effective than training directly on the exam task due to gains of motivation. The work on deep and shallow learning suggests this is possible.

Problem-based (project) work: If you want to teach an ability to do projects, then obviously the principle of requiring the target activity will require learners to do project work. If not, then the question is whether motivational and possibly understanding gains will nevertheless make it better than direct training on the target (examined) task.

References

- Bartlett, F.C. (1932) "Remembering" (Cambridge, U.K.: Cambridge University Press).
- Berlin, B. & Kay, P. (1969) Basic color terms: their universality and evolution (U. California Press)
- Byrne, M., Johnstone, A.H., & Pope, A. (1994) "Reasoning in science: a language problem revealed?" School science review vol.75 no.272 pp.103-107
- Cassels, J.R.T. & Johnstone, A.H. (1985) Words that matter in science: a report of a research exercise (London: Royal Society of Chemistry)
- Cole, M. & Scribner, S. (1977) "Cross-cultural studies of memory and cognition" ch.8 pp.239-271 in R.V.Kail & J.W.Hagan (eds.) Perspectives on the development of memory and cognition (Erlbaum: Hillsdale, N.J.)
- Edwards, D. & Mercer, N. (1987) Common Knowledge: the development of understanding in the classroom (Routledge: London).
- Glaserfeld — see von Glaserfeld, E. (1984)
- Goodman, N. (1979) Fact, fiction and forecast (Hassocks, Sussex: Harvester Press)
- Jonassen, D.H. (1991) "Objectivism vs. constructivism: do we need a new philosophical paradigm?" Educational technology: research and development vol.39 no.3 pp.5-14
- Katona, G. (1940/1967) Organizing and memorizing: studies in the psychology of learning and teaching (New York: Hafner).
- Kripke, S.A. (1982) Wittgenstein on rules and private language (Blackwell: Oxford)
- Laurillard, D. (1993) Rethinking university teaching: A framework for the effective use of educational technology (Routledge: London).
- Marsh, H.W. (1984) "Experimental manipulations of university motivation and their effect on examination performance" British journal of educational psychology vol.54 pp.206-213.
- Marsh, H.W. (1987) "Student's evaluations of university teaching: research findings, methodological issues, and directions for future research" Int. journal of educational research vol.11 no.3 pp.253-388.
- Marton, F., D.Hounsell & N.Entwistle (eds) (1984) The experience of learning (Scottish academic)
- Newell, A. (1990) Unified theories of cognition (Harvard University press: London)
- Perry, W.G. (1970) Forms of intellectual and ethical development in the college years (New York: Holt, Rhinehart and Winston)
- Perry, W.G. (1988) "Different worlds in the same classroom" in P.Ramsden Improving learning: new perspectives (London: Kogan Page)
- Piaget, J. (1970/1983) "Piaget's theory"
 a) Carmichael's manual of child psychology (1970)
 b) P.H.Mussen (ed.) Handbook of child psychology 4th edition (1983); W.Kessen (ed.) vol.1 History, theory, and methods ch.3 pp.103-128
- Putnam, H. (1975) The meaning of meaning, in Mind, language and reality, (CUP, Cambridge)
- Quine (1969) Preface to D.K.Lewis Convention: a philosophical study (Harvard UP)
- Rumelhart, D.E. & Norman, D.A. (1978) "Accretion, tuning and restructuring: three modes of learning" in J.W.Cotton & R.L.Klatzky (eds.) Semantic factors in cognition pp.37-53 (LEA)
- Scribner, S. & Cole, M. (1973) "The cognitive consequences of formal and informal education" Science

vol.182 pp.553-559

Svensson, L. (1984) "Skill in learning" in Marton,F., D.Hounsell & N.Entwistle (eds.) The experience of learning (Edinburgh: Scottish academic press) ch.4 pp.56-70

von Glasersfeld,E. (1984) "An introduction to radical constructivism" in P.Watzlawick (1984)

Watzlawick, P. (1984) (ed.) The invented reality: How do we know what we believe we know? Contributions to constructivism (W.W.Norton: New York)

Wittgenstein, L. (1953) Philosophical Investigations (Oxford: Blackwell)