

LANGUAGE TEACHING IS TEACHING LANGUAGE¹

‘Something must be done. This is something. Therefore let us do it.’ (*Yes, Minister*)

knowledge and skills; prioritising

I should like to start by putting forward a rather unoriginal view of the task facing us as language teachers, teacher trainers or course designers. This is that:

- Language use involves: 1) a knowledge base, and 2) skill in performing operations which draw on this knowledge base. These are equally important.
- Compared with, say, music or driving, the knowledge base required for effective language use is vast.
- Both the knowledge base and the associated skills take a great deal of time and (for most adults) considerable work to assimilate.
- Most learners only have time to master a small part of a foreign language.
- Our task is therefore 1) to prioritise, selecting the language and skills which are most important for our learners, and 2) to ensure that our learners engage with the language and skills selected in ways which will ensure that they are effectively learnt.

Most people in ELT are, of course, well aware of these considerations, and plenty of very effective teaching goes on all over the world as a result – I don’t wish to suggest that everybody is getting it all wrong. However, there does seem to be a perennial tendency for the balance to tip over too far to one side or the other of the knowledge/skills divide; and also for means – teaching activities – to supplant ends – the knowledge and skills that the activities are supposed to be teaching. I shall argue that this is happening now, and indeed that, in some corners of the profession at least, we may be so concerned with the teaching of skills – language as process – that we are in danger of seriously neglecting the knowledge base – language as product.

doing things and teaching things

When I was taking my first steps in EFL, language teaching seemed a relatively simple business. We taught – and we took it for granted that students learnt – grammar, vocabulary and pronunciation. The fluent deployment of these elements was practised through work on receptive and productive skills. (It was well known that there were four of these.) Everything was neatly packaged in the textbook. There were, it is true, occasional hints that we might be overlooking something (as when students asked us to ‘teach them conversation’ and we shuffled our feet and told them uneasily that conversational fluency would just ‘come with practice’.) But by and large the system worked smoothly: we gave our lessons and students’ English got better.

Something that worried me even in those early days, though, was a feeling that we tended, without realising it, to slide from teaching things into doing things. This happened most obviously after one left the simple certainties of the elementary

¹ Written-up version of plenary delivered at IATEFL annual conference 1996, published in *IATEFL Annual Conference Report 1996*.

syllabus. If you stopped a beginners' teacher on the way to the classroom and said 'what's going to happen in your lesson?', you were likely to get an answer in terms of goals: 'colour words', 'the present perfect', 'things to say in shops'. But if you tried the same thing with an intermediate teacher, the answer was much more likely to describe activities: 'dictation', 'reading', 'doing a dialogue', 'making a radio programme' While one told you what she was going to teach, the other said what she was going to do.

In itself, this change of emphasis is quite understandable. As students become more proficient, their needs become more varied and diffuse, and it becomes harder to define one's aims in terms of a list of specific linguistic products. At the same time, skills practice inevitably takes on increasing importance – intermediate students typically know a lot of language which they can't use, and more advanced students may need training in special uses of language relevant to their professional needs. So it is natural that, from intermediate level upwards, types of activity become important elements in teaching syllabuses.

But the change of emphasis brings with it a serious danger of losing focus on goals. The activities we select can become ends in themselves, while the language they are supposed to be teaching gets pushed into the background. We can easily end up simply doing a lot of things which seem vaguely related to language teaching, keeping students happily occupied with tasks that involve enough interactive use of language to reassure us and them that we are doing our jobs. (Surely, we feel, if they're speaking/reading/writing/hearing English they must be learning it?) So while we may intend, in theory, to give equal importance to the knowledge base and the associated skills, our students may actually spend a great deal of their time doing rather ill-defined fluency practice, and not very much time systematically learning new language.

everything is getting more complicated

Since my early days in the profession, language teaching has progressed enormously. We know much more about language, we have more sophisticated ideas about how people learn it, and we have far better ways of teaching it. Work done by researchers in discourse analysis, in particular, has made it possible to tabulate the ways in which real-life exchanges work, to discover how different kinds of texts are structured, and to describe language and its use in terms of semantic categories such as 'notions' and 'functions'. (So if students ask us now to 'teach them conversation', we know what the request means and we have a good idea of how to go about meeting it.)

Unfortunately, as we learn more about language, and as our methodology develops to keep up with our knowledge, it becomes even harder to define clear aims and to choose appropriate activities by which to achieve them. There is so much to teach, and so many things to do. Teachers, trainers and materials writers who wish to be properly informed have to battle their way through a dense jungle of facts and theory, often impeded by thickets of vicious terminology. And with all of this, we don't seem much nearer to answering the central question: 'What happens in people's heads when they learn languages, and how can we make it happen more effectively?' In moments of gloom I have sometimes been reminded of Mark Twain's words: 'The researches of many commentators have already thrown much darkness on this subject, and it is probable that if this continues, we shall soon know nothing at all about it'. As we despair of ever getting a clear overview that would enable us to make rational

decisions, we are naturally tempted to stop trying and unconsciously take refuge in a cocktail of language-practice activities.

language as process

The perennial tendency to seek refuge in user-friendly activity has received a powerful boost from one current in contemporary linguistics which has had a good deal of influence on language-teaching – the move towards seeing language as ‘process’ rather than ‘product’. This has helped to legitimise a skills-centered view of what we are doing: our job (as many people now see it) is not to supply the learner with a product (the language); it is to enable him or her to engage in processes which will inculcate the skills he or she needs for successful language use (or some similar formulation). Many current language courses are based on what one might call the ‘battery of skills’ approach, in which the principal focus is not so much on teaching language as on training people to do things with it. When I started teaching, reading was a skill. Fifteen years later, in one well-known analysis, reading was alleged to comprise nineteen separate skills. Inventories of skills may be supplemented by equally complex inventories of strategies – learning strategies, communication strategies and so forth – and it is often suggested that these too should be taught.

‘Although ‘strategy’ has never been adequately defined in the learner language literature, and although some bizarre labels are given to learner behaviour, such as ‘the strategy of incorrect application’, it has been widely exemplified, and it comes over as an altogether positive concept: learners deploying strategies or teachers encouraging learners to use their existing strategies and add more to their strategy repertoires, seem assured of receiving an accolade.’ (James 1990)

This is not the place for a detailed criticism of the ‘skills and strategies’ view of language, but I should at least like to suggest, as a useful operating principle, that we should avoid teaching any ‘skill’ or ‘strategy’ unless we are reasonably certain (1) that it really exists, (2) that the learner needs it, (3) that the learner does not already possess it and (4) that it can actually be both taught and learnt.

A process view of language, then, with its skills-centered view of aims as well as methodology, moves us further in the direction of doing things rather than teaching things. Classroom activity, in this view, is now no longer a means to an end which includes the acquisition of the knowledge base; it can become virtually an end in itself, with the accompanying risk that the language itself gets swept even further under the carpet.

simplification: the search for a quick fix

With so many things to teach, and so many things to do, it is natural that teachers and course designers should look for ways of simplifying their task. At the same time, researchers are under considerable pressure to come up with practical solutions to the language-learning problem. An enormous amount of time and money is spent worldwide teaching languages to children and adults; in general, the results are not spectacular. Anybody who can make this massive investment more cost-effective by finding ways of speeding up language-learning will earn the gratitude of nations and achieve fame and funding – or at least, in these hard times, have a better chance of surviving the next round of staff cuts. It is not surprising, then, that new ‘methods’ and ‘approaches’ appear at regular intervals. Up to a point this is desirable and indeed necessary: at the very least, new approaches generate new energy, and often they

result in real, if modest, progress. The danger is that in our search for greater efficiency, and for simple routes through the conceptual jungle of language-teaching, we may latch onto the latest development and turn it into a ‘quick fix’ that will solve all our problems, creating one homogeneous ‘method’ or ‘approach’ that, at last, we can use to teach the language twice or three times as effectively as before. (I was recently shown a description of a well-known ‘method’, by no means new, that is alleged to teach languages up to ‘six times as fast’. I was discouraged from trying it out by the reflection that if somebody had invented a fuel, twenty years ago, that was six times as efficient as petrol, nobody would be trying to convert us to it today – we would all be using it.)

Looking back over the last few decades, it is disturbing to realise how many methods, approaches and technologies have been perceived as ‘the key’: structure drilling, the language lab, the audiovisual approach, the input hypothesis, the silent way, humanistic approaches such as counselling learning, suggestopaedia, the use of authentic materials, the notional/functional syllabus, the communicative approach, total physical response, learner-centered approaches, task-based syllabuses – to name but a few. In the wake of the seventies (the heyday of gurus and miracle methods), Catherine Walter and I published an article in *ELT Journal* on ‘Teaching English by Sensory Deprivation’. Such was the climate of the times that more than one reader took the paper seriously, and references to it started turning up in bibliographies.

I think it is probably wise to beware of any philosophy of language teaching that has a name: ‘the X method’ or ‘The Y approach’. The very fact that the method is delimited in this way means that its proponents are focusing primarily or exclusively on one aspect of language; inevitably, therefore, other aspects will be neglected. New approaches and technologies are good servants (because of what they add to our professional repertoire), but generally bad masters (because of what they make us leave out). When I was first learning my trade, we were good at teaching grammar but bad at teaching conversation. More recently, the opposite has often been the case. A few weeks ago somebody told me ‘We’ve all gone over to the Lexical Approach now – we hardly do any grammar at all’. I doubt if Michael Lewis would have been pleased to hear that his ideas had been allowed to fill somebody’s whole horizon in this way.

As each new panacea turns out not, after all, to be the miracle cure, disillusionment sets in, and another pendulum swing starts; so that too much of our energy goes into reacting to our mistakes, and not enough into simply trying to achieve our basic language-teaching aims. As Ian Stewart put it in *New Scientist*, talking about a topic remote from language teaching: ‘It has been known for a long time that control systems can behave chaotically ... if they are ‘over-driven’ ... the whole system thrashes, spending nearly all of its time reacting to its own errors and very little time reacting to the reality it is supposed to be controlling’. (Stewart 1995)

‘the language will take care of itself’: task-based syllabuses

Neglect of the knowledge base is not always an accidental side-effect of a concern to teach skills; it can be a matter of deliberate policy. Prabhu (1987) refers to ‘a strongly-felt pedagogic intuition that the development of competence in a second language requires *not* systematisation of language inputs or maximisation of planned practice, but rather the creation of conditions in which learners engage in an effort to cope with communication’. Many people working in this perspective hold the view that teaching and learning can be organised according to a purely task-based syllabus. According to

some versions of this model, if one analyses the ‘target’ tasks that learners will have to perform in real life, and if one then takes learners through examples of similar tasks in the classroom, the language (often described, significantly, as ‘skills’) that is needed will be generated in the process: it will emerge naturally from the materials used for the task, or from the teacher’s input, or from the classroom interaction. There is therefore no need to draw up explicit formal or functional language syllabuses: if the tasks are correctly chosen, the language will take care of itself.

Well, will it? Will a specification of behavioural objectives automatically generate, as an incidental payoff, a specification of the knowledge base that learners will need for an adequate command of the language? If you take learners through a set of carefully selected tasks, will all the high-priority grammar, pronunciation and vocabulary points show up and be learnt?

Let us make the precarious assumption that all the grammar that students are likely to need in real life will indeed be thrown up, without special planning, by a course consisting only of tasks; and that this grammar will be learnt effectively either during the task-based activities or during additional follow-up ‘form focus’ sessions. Let us also, for the sake of argument, accept Tom MacArthur’s memorable definition of pronunciation as ‘that part of a student which is the same at the end of a language course as at the beginning’. That leaves vocabulary.

In discussions of task-based learning, vocabulary tends to be invisible. (In two recent books by Nunan (1988, 1989), there are no index references to either ‘vocabulary’ or ‘lexis’.) This is strange, because it seems obvious that, at least after elementary level, the largest part of a language learner’s task is to build up an adequate stock of high-priority words and lexicalised phrases, including both knowledge of their forms and an awareness of their more important meanings, the major collocational and syntactic constraints on their use, and so on. If a purely task-based syllabus is to do its job, it should then provide learners not only with the grammar, but also with the vocabulary they need. Unfortunately, it seems clear to me that there is no way in which it can do this effectively: any task-based syllabus must be supplemented, I believe, at least by a separately-planned lexical syllabus.

the need for a lexical syllabus

In order to see why separate lexical syllabuses are necessary, we must consider questions of extent and frequency. Let us take the case of a typical learner – call her Sophia. Sophia is a reasonably well-motivated intermediate student, moving into her third year of English in a secondary school in a non-English-speaking country where task-based language-teaching syllabuses have been introduced. She has already learnt the commonest two thousand-odd vocabulary items, and she would like to finish her third year knowing the next thousand or so high-priority words and expressions. Now some of the vocabulary items in this frequency-range will be ‘task-bound’, and are sure to be learnt if the relevant tasks form part of the syllabus – for example, if Sophia and her fellow-students work on business negotiations, then one does not need to worry about whether they will learn *contract*, *meeting*, *negotiate*, *confirm*, *delivery date* and so on, because these are virtually certain to come up in one way or another. Unfortunately, however, the bulk of intermediate vocabulary is not task-bound in this way: there is a great deal of what you might call ‘general-purpose’ material. The ‘third thousand’ items include, for example, *calm*, *noisy*, *swallow*, *take trouble*, *rubbish*, *lane*, *genuine* and *out of sight* – but it is hard to see what tasks could be guaranteed automatically to throw up these words and expressions.

Perhaps, though, even the general-purpose items that Sophia needs are frequent enough to come up automatically in her year's work, whatever the tasks covered. In order to see whether this is a reasonable supposition, we need a little mathematics. During the school year, assuming a fair amount of written and recorded material, plenty of classroom activity, a teacher who talks a lot and some reading outside class, Sophia might be exposed – on a very generous estimate – to around 350,000 words of input: 10,000 words a week. Now words like *calm*, *noisy*, *swallow*, *rubbish*, *lane* and *genuine* tend to have frequencies between 7 and 30 occurrences per million running words. This means that any one of them, taken singly, is very likely to come up at least once in a corpus of 350,000 words. Using the frequency figures from the LOB Corpus, and applying some probability theory, *calm* would have better than a 99.9% chance of occurring; *noisy* a 95% chance; *swallow* a 90% chance. However, things that each have a good chance of happening separately have a far less good chance of all happening together (this is why one or two of your numbers often come up in the lottery, but never all six). The chance of twenty-five specific words at this level of frequency all coming up in our hypothetical year's work is in fact not much better than evens – 50% or so. Assuming that 600 of our 'third thousand' high-priority items are general-purpose vocabulary, not automatically thrown up by the tasks chosen, then the chance of their all coming up anyway, without being artificially fed in, is quite remote – of the order of a million to one against. Sophia will not learn her thousand words.

Vocabulary, then, will not take care of itself. If students with limited time available for study are to learn high-priority lexis, this needs to be deliberately selected and incorporated into learning materials or activities. If this is not done, students will not be exposed – even once – to numerous important vocabulary items, and they will finish their courses with serious gaps in their knowledge. In earlier times we provided students, so to speak, with the necessary bricks, tiles, timbers, mortar and so on, and assumed that they could build their own houses. Exclusively task-based approaches fall into the opposite error: they get learners building houses right away, but assume that the various supplies needed will magically materialise – as if delivered by elves – at the right times and in the right quantities. Unfortunately elves are scarce in ELT: we need lexical syllabuses.

Unlike the hypothetical third-year student in our example, some learners follow language courses in the country where the target language is spoken, and here the issues are of course somewhat different. The massive exposure to the language which such learners can get outside the classroom will guarantee that they acquire high-priority vocabulary, and class time can be used for other things. Whether purely task-based approaches are any more suitable for this environment is another question: on the face of it, it would seem that outside exposure will also equip students, on the whole, to do the things in English that they need to learn to do. One is led to ask what a language course in the target-language country is actually for; and this is not, perhaps, an easy question to answer.

learner independence

'Language teaching and learning have frequently been beset by techniques in which the tail wags the dog. Thus the language laboratory – essentially a useful technique – became a controlling factor in some methodologies, so that language learning and teaching was organised around the language laboratory. Self-instruction and self-directed learning may pose a similar threat of a

particular learning mode taking over the whole of the learning programme and distorting it so that the covert aim becomes the success of self-directed learning rather than the successful learning of the target language.’ (Dickinson 1987)

Current resistance to pre-planned language syllabuses often arises from an understandable desire not to impose external constraints on the very personal business of language learning. Each learner is unique, with his or her own goals, learning styles, and so on, and it is natural to feel that the more we can involve learners in the decisions that affect their learning, from personalisation of exercises right up to overall syllabus design, the more chance we have of providing courses that are sensitive to their needs.

It is of course true that only the learner really knows exactly what he or she wants. It is, however, equally true that only the teacher knows what there is to be learnt. In other branches of teaching we are not usually so diffident about imposing direction and constraints on the learners – I would not, for instance, encourage my seventeen-year-old son to adopt a discovery approach to learning to drive my car; nor would I want my rock-climbing instructor to leave me to find out for myself how to rappel down effectively.

It is not altogether clear, in fact, that all learners want the degree of autonomy that some teachers think is good for them. In two Australian studies investigating learning preferences quoted by Nunan (1988), while teachers gave most importance to conversation practice, self-discovery of errors and pairwork, learners gave higher ratings to pronunciation practice, explanations by the teacher, error correction and vocabulary development. There was, indeed, some evidence of irritation with activity-rich approaches: as one student expressed it, ‘I don’t want to clap and sing, I want to learn English.’

natural acquisition versus instruction

Recent years have seen a good deal of research on first and second language acquisition, and on the similarities and differences between them. This has led many researchers to ask whether ‘instructed learning’ can really achieve the same kind of results as ‘natural acquisition’. A well-known extreme view is Krashen’s ‘no interface’ position: that conscious learning of grammar rules provides learners, at best, with a kind of knowledge that can be used to monitor their own production when there is time to do so, and that this knowledge cannot be used to generate spontaneous grammatically-accurate utterances. Subsequent studies have tended to rehabilitate grammatical instruction up to a point, showing that rule-learning can have at least a modest effect on accuracy (see Ellis 1994 for discussion). However, there is still quite a widespread feeling that rule-learning is ‘artificial’, and that the more we can approximate the conditions of natural acquisition in our teaching, the better we are likely to do.

Strangely, the ‘instruction versus natural acquisition’ debate has concentrated almost exclusively on the learning of grammar, and on the question of *how* grammar is learnt. In fact, it seems to me that this question is almost irrelevant to a comparison between instructed learning and natural acquisition. The crucial difference between the two, surely, has to do with vocabulary, and far more with *what* is learnt than with *how* it is learnt.

According to a study carried out in the 1970s (Carey 1978), English-speaking children have learnt on average 14,000 words by the age of 6. This works out at about 6.3 words a day. Adults learning a second language at this rate would need around

10,000 hours' exposure to acquire a vocabulary of 5,000 words – a fairly time-consuming way to reach an intermediate level. Unaided natural acquisition, then, works slowly; and, as we have already seen, teaching approaches which simulate the conditions of natural acquisition cannot even present learners with the vocabulary they need in the time normally available for language courses.

The point about instruction is that, properly organised, it can do two things. First of all, it can *select* the language which will be presented to the learner, making sure that, in the short time available for a language course, as many high-priority items as possible occur in the input. Secondly, it can supply *concentrated exposure to*, and *concentrated practice of* the items presented, by providing activities which force the learner to engage intensively with the new material. Whether or not this is effective depends on too many factors to list, but at least planned instruction provides learners with some chance of acquiring essential vocabulary; whereas an approach which does not even get this vocabulary into the input simply abdicates responsibility for teaching core aspects of the language.

summary and conclusion

I would, then, like to see something of a 'return to basics' in our profession (without wishing the expression to carry the obscurantist overtones that it has recently acquired in the context of political crowd-pleasing). Despite the enormous progress that has been made in language analysis, syllabus design and methodology, it seems to me that our teaching aims are not actually very different from those that were identified when I first went into the classroom. That is to say, our task is to provide learners with a command of selected high-priority aspects of grammar, vocabulary and pronunciation, and with facility in using these accurately and appropriately. (Though of course we do not mean quite the same by words like *grammar*, *vocabulary* or *appropriately* as we did thirty years ago.) This return to basics would involve, among other things:

- A rehabilitation of instruction, in the quarters where this is needed, together with a clear understanding of the need to select and present input, especially lexis, in a principled way, and of the consequent limitations of learner autonomy and of approaches that emulate 'natural acquisition'.
- The realisation that a trained language teacher needs a thorough knowledge of the structure of the language that he or she is teaching – native-speaking teachers of English today are often surprisingly ignorant of grammar.
- An acceptance that language learning is hard, and sometimes dull, and that it will remain so – there are no quick fixes. Of course we must find ways of making it interesting, but we must remember that what the bored teacher wants is not necessarily what the student needs. We naturally doesn't want to be doing the same things in our twentieth year of teaching as in our first; but our learners' needs may not be very different today from those of their parents. And if in our first year we were teaching vocabulary efficiently, and in our twentieth year we are training learners in autonomous interactive discourse negotiation strategies, we need to ask if this really represents progress, and if so, for whom.
- A sharpened focus on product as well as process. There are an enormous number of good things to do in language lessons, but we do need to be sure that the activities we choose actually teach something useful. 'Product' is not a dirty word – it is what our learners are paying for. When we walk out of a classroom, if the lesson was successful, we ought to be able to say what important bit of language

the learners now know that they didn't know before; or what important skill they can now operate a little better than they could before. If we can't do this, there may be something wrong.

- A rejection of the view that communication is an absolute good. The more we can integrate real communication into language practice the better, but it does not follow that because students are communicating they are learning English; and some activities (such as learning by heart or mechanical structure-practice), unfashionable because they are totally uncommunicative, may none the less be very valuable.
- A common-sense attitude to what we are doing. It can be helpful, in evaluating a fashionable approach, to ask how we would like it applied to ourselves. If we were learning Chinese, Greek or Swahili, how many of our own methodological assumptions would we put up with? (Not all that long ago, teachers in Britain regularly forbade their students to translate or use bilingual dictionaries. Guess what kind of dictionaries they all took on holiday abroad.)

Kenneth Tynan, reflecting on his own profession, said that a good drama critic is one who perceives what is happening in the theatre of his time, but that a great drama critic is one who perceives what is not happening. Our conference programmes, journal articles and publishers' catalogues show evidence of a great variety of concerns, pursued with enormous creativity and impressive energy. As we find ever more interesting, motivating and professionally rewarding things to do, we, too, need to look carefully at what is not happening.

References

- Carey, S.** 1978. 'The child as word learner' in Halle, M., Bresnan, J. and Miller, B. (eds): *Linguistic Theory and Psychological Reality*. Cambridge, Mass: MIT Press.
- Dickinson, L.** 1987. *Self-instruction in Language Learning*. Cambridge: Cambridge University Press.
- Ellis, R.** 1994. *The Study of Second Language Acquisition*. Oxford: Oxford University Press.
- James, C.** 1990. 'Learner language.' *Language Teaching* 23.4. Cambridge: Cambridge University Press.
- Nunan, D.** 1988. *The Learner-Centered Curriculum*. Cambridge: Cambridge University Press.
- Nunan, D.** 1989. *Designing Tasks for the Communicative Classroom*. Cambridge: Cambridge University Press.
- Prabhu, N.S.** 1987. *Second Language Pedagogy: a Perspective*. Oxford: Oxford University Press.
- Stewart, I.** 1995. 'Juggling by Numbers.' *New Scientist* 20 May 1995.