
Evaluation of Learning Technology Implementation

Professor Barry Jackson, Middlesex University

Abstract

Learning technologies present the opportunity to augment or replace the role which was traditionally played by the teacher. In a learning-led educational context the effectiveness of teaching, whether by the teacher, or through interaction with media, can be readily evaluated by the degree to which it contributes to learning. The use of level descriptors, such as Biggs' SOLO taxonomy, provides a powerful tool to assess the outcomes of learning, and by implication the effectiveness of the learning activity. Contextual factors which correlated with the SOLO levels, and encourage deep approaches to learning, can be used to improve effectiveness. This paper outlines the SOLO taxonomy, indicates implications for implementation, and discusses the correlating factors which enhance students' approach to learning.

Purposes and scope of evaluation

'Evaluation is any activity that throughout the planning and delivery of innovative programmes enables those involved to learn and make judgements about the starting assumptions, implementation processes and outcomes of the innovation concerned'
Stern [1988]

Under this definition the principal purpose of evaluating learning technologies is to provide the designer or user with enough evidence on which to make confident judgements regarding the effectiveness of the innovation. From these judgements actions can be based which will result in the redesign or adaptation of the implementation to improve its performance. Judgements of the 'effectiveness of the innovation' are in effect judgements of the extent to which the innovation meets its identified aims - its fitness for purpose. It is important therefore that these intentions are thought through and clarified at the design stage of the project - good evaluation practice is designed-in, not appended, to projects.

There are three stages in innovative projects about which judgements may be made: intentions,

implementation and outcomes. The match between the outcomes and intentions provides a measure of the success of the innovation. If outcomes do not match the intentions, then it is readily observed that something needs to be done: perhaps implementation is inappropriate, or perhaps the starting assumptions should be revisited and revised. Even when outcomes closely match the intentions, and the project is counted successful there may be assumptions, behind the stated aims of the innovation, often implicit and unstated. An innovation may be judged as 'fit for purpose' while the fitness of the purpose itself is questionable. It is important therefore to be as explicit as possible about the starting assumptions, in order that the evaluation of design and implementation can be most useful.

To take more concrete examples: a learning package might be evaluated against a number of intended outcomes.

These might include:

- ease of use - can the target users navigate and interact with the product in the expected way?
- efficiency - does the product provide efficient coverage of the appropriate curricular unit? Can more multiple learners access the content than by other means?
- preferences - do learners enjoy it? will they choose to use it if it is offered? to what extent does it cater for variety of user?
- attractiveness - is it sufficiently attractive to persuade purchasers/tutors to choose it over alternatives
- cost effectiveness - is it cost effective in comparison with alternative means of achieving the same intended outcomes?

However, at another level there are assumptions about the nature of learning and learners, which will shape and define the answers to all these evaluative questions. It is to these assumptions that I wish to draw attention below - assumptions about learning at tertiary level which I believe are shared by most teachers, and which provide a necessary aspect of evaluation of any learning technology, alongside any of the aspects outlined above.

Evaluation of learning

The evaluative question I propose is simply stated thus: does the product provide opportunities for learning at an appropriate level? To answer this apparently simple question it is necessary to explore what we mean by 'learning' and 'appropriate level', since these are the assumptions which will shape the design of teaching materials, and their evaluation.

There are broadly two traditions of theories applying to higher education. The first is the objectivist tradition, which locates knowledge outside the knower. In this tradition knowledge exists independently of the knower, and understanding is coming to know that which already exists. In this tradition, knowledge is seen as independent of particular contexts: teaching is a matter of transmitting this knowledge and learning is a matter of receiving, storing and applying it.

The second tradition rejects the separation of knower and knowledge. In this, the constructivist tradition, knowledge is seen to be constructed by the knower in acts of understanding. Meaning is created by the learner, not simply received. Learning is context-sensitive: the learner brings to any learning situation an accumulation of assumptions, motivations, conceptions and previous knowledge, which will largely determine the nature and quality of learning which takes place. David Boud succinctly draws our attention to the central importance of the learner's acts while learning:

'ultimately it is only the decisions which learners make about what they will or will not do which actually influence the outcomes of their learning.'

[Boud, 1981]

Either of these broad theoretical positions may inform the the teaching decisions made by individual teachers, or indeed the design decisions made by designers of learning materials or activities. Argyris [1976] drew attention to the difference between espoused theories, that are expressed as the theories underlying practice, and theories-in-use, which are the actual, unexpressed theories which guide practice in reality. Espoused theories and theories-in-use are often different. Although constructivism is now the dominant espoused theory in higher education, it is possible to see in current practice, including the design of many learning activities using C&IT, that objectivism is still the dominant theory-in-use. Biggs [1996] reminds us that professionalism

requires the espoused theory to be the theory-in-use.

For the purposes of evaluation it is important to understand that the stated intentions underlying the innovation may not reflect the actual theory-in-use, and that a product which aims to develop students' understanding might have been designed with an objectivist approach. Such a product is unlikely to be successful in achieving appropriate aims.

The commonly stated aims of learning in higher education are closely associated with a constructivist approach. That is to say that what is said to be valuable in higher education is something more than the simple acquisition of skills or accumulation of knowledge. Academic learning at tertiary level is believed to be fundamentally about the development of understanding and the ability to apply critical judgements to presented knowledge. It aspires to the position in which students construct their own maps and networks of meaning, testing them against principles and descriptions by others. Particular skills and elements of knowledge might play significant parts in the construction of this personal understanding, and will need to be mastered by the learner, but these will be means and not ends in themselves. And such skills may be better achieved when their development is closely integrated with associated conceptual learning. (As an example, the reason that we wish science graduates to be skilled in laboratory techniques may not be so much that those skills will be of applied value in later life, but rather that by gaining the skills the learners are enabled to further develop their knowledge and understanding).

The implications of this are that learning technologies, like other teaching situations, should be evaluated as to the degree to which they encourage or facilitate the development of understanding. The exceptions to this might be learning programmes whose sole aim is skill development.

The development of understanding is a complex process, and its outcomes have been typically difficult to identify and classify. However, the work of Biggs and Collis [1982] has provided a powerful taxonomy, The Structure of the Observed Learning Outcome or SOLO taxonomy, and this might provide a useful tool to assist the evaluation of learning technologies.

SOLO and levels of learning outcome

The SOLO taxonomy provides a systematic way of describing a hierarchy of complexity which learners

show in mastery of academic work. The taxonomy, arrived at through phenomenographic research, is intended to be and has been successfully applied as, a description of the range of performances produced by learners in attempting a particular academic activity, and its particular strength is its generality - that is it is not content dependent, and it may be used effectively across a number of subject areas (see for example case studies in Gibbs [1993]).

SOLO describes five levels of sophistication which can be encountered in learners' responses to academic tasks:

- Prestructural - the task is not attacked appropriately, the student hasn't understood the point;
- Unistructural - one or a few aspects of the task are picked up and used (understanding as nominal);
- Multi-structural - several aspects of the task are learned but are treated separately (understanding as knowing about);
- Relational - the components are integrated into a coherent whole, with each part contributing to the overall meaning (understanding as appreciating relationships);
- Extended abstract - the integrated whole at the relational level is reconceptualised at a higher level of abstraction, which enables generalisation to a new topic or area, or is turned reflexively on oneself (understanding as far transfer and as involving metacognition).

[Biggs 1996]

Levels 4 and 5 can be seen to be qualitatively different from the lower levels, in that both 4 and 5 involve the learner in integrating the new knowledge and skills into a coherent structure. The learner is making meaning. It is learning at this level which is characteristic of effective learning in higher education, and is the desired aim of most established programmes.

This provides a possible approach to evaluating the effectiveness of learning and teaching situations for facilitating students' understanding: if evidence of higher SOLO levels can be found then it suggests that the learning activity is effective at encouraging the construction of knowledge. Absence of evidence might suggest that something in the design or implementation of the learning activity is in need of improvement.

There are two important problems to be overcome with this approach. Firstly there needs to be an effective way of identifying the different SOLO

levels of outcome which students achieve. It is not enough to rely on assessment results to indicate these, since that presupposes that the assessment scheme is designed to measure, or actually does measure, the level of understanding. This is by no means common (nor unproblematic). It is unlikely that assessment by multiple choice questions is effectively measuring the level of understanding, for example.

The most fruitful route to gathering this data is by analysis of students' written work, using a protocol based on the SOLO taxonomy. With experience and sensitivity this can provide evidence for reasonably robust judgements about the level of understanding being achieved, independently of the content. Alternatively analysis can be made of students' reflective written or verbal reports of their learning activities, elicited for example by interview (eg case studies in Gibbs [1993]) or in learning journals. Biggs [1996] provides a starting point for the development of an evaluative protocol, focussing on the *effect* of understanding: if you understand something properly you act differently in contexts involving the content understood. Biggs proposes a hierarchical list of 'performances of understanding', from most desirable to barely satisfactory using SOLO as a baseline. As a performance measure the list focusses on verbs. An example based on a particular unit in a BEd programme at University of Hong Kong provides this example of a descriptor for the Most Desirable (extended abstract) performance:

'metacognitive understanding, students able to use the taught content in order to reflect on their own teaching, *evaluate* their decisions made in the classroom in terms of theory, and thereby *improve* their decision-making and practice. Other outcomes: *formulating* a personal theory of teaching that demonstrably drives decision-making and practice, *generating* new approaches to teaching on the basis of taught principles and content'.

Biggs [1996]

It is interesting to compare this with a descriptor for level 3 Moderately satisfactory (multistructural):

'students understand declaratively, in that they can *discuss* content meaningfully, they *know about* a reasonable amount of content, but don't transfer or apply it easily'.

Biggs [1996]

In summary, a performative picture of

understanding offers a way in which teachers and designers of learning technologies can develop means of judging the level of students' learning in particular contexts.

The second problem to be addressed is a larger one; so large in fact that it cannot be adequately treated within the scope of this paper, beyond drawing attention to it. A constructivist view of learning, as noted above, recognises the central importance of the learner and the significance of the context on learning. Consequently any evaluation of particular teaching methods or learning technologies is fraught with difficulty - the effectiveness will vary, depending on the scope of previous knowledge, attitudes, and conceptions which particular learners bring to the learning situation, and the larger context in which the learning situation is embedded. The outcomes of a particular evaluation of learning effectiveness cannot therefore be easily transferred out of the context in which the evaluation occurs. Evaluation of learning technology must be situated in the context in which the technology is used. In effect this is to say that evaluation of *implementation* is unavoidably needed, and that the judgements made by evaluators may say less about the design and more about the implementation of the technology in question. And this is entirely appropriate, of course. Effective learning has often arisen from a bad book introduced at the right time.

References:

- Argyris, C.**, [1976] *Theories of action that inhibit individual learning*, American Psychologist 31: 638-654
- Biggs, J.**, [1996] *Enhancing teaching through constructive alignment*, Higher Education 32: 347-364
- Biggs, J.B., and Collis, K.F.** [1982] *Evaluating the quality of learning: the SOLO Taxonomy*, New York: Academic Press
- Boud, D.**, [1981] *Developing student autonomy in learning*, London, Kogan Page
- Gibbs, G.** [1993] *Improving the quality of student learning*, Bristol, TES
- Stern, E.**, [1990] *The evaluation of policy and the politics of evaluation*, in The Tavistock Institute of Human Relations Annual Review