Approaches to evaluating the effect of ICT on student learning

Rob Phillips and Tony Gilding

Introduction

This starter guide gives a broad overview of issues surrounding the evaluation of the effectiveness of Information and Communications Technology (ICT) environments. Unfortunately, there is no 'magic bullet' which will tell you simply and easily that a certain application of ICT has had an effect on student learning. Indeed, we can't evaluate ICT in isolation – we have to look at student learning in the teaching and learning environment as a whole.

This guide is aimed at teaching staff and project managers involved in developing ICT applications. We will focus on the hardest part – planning an evaluation study. Once a plan is in place, people with social science research skills can be employed to do much of the leg-work of the study.

What is meant by evaluation?A useful broad definition of evaluation is "Providing information to make decisions about the product or process*".* However, when we try to apply this definition to "evaluating the effect of ICT on student learning", we run into problems. What is the product? What is the process? These are questions that are difficult to answer when applied to learners. Because human beings are complex creatures, and learning is a complex, multifaceted activity, we are led from evaluation into research.

While we might evaluate the usefulness of ICT in an educational context, we also need to research how ICT can affect the processes of learning, and what learning outcomes are achieved. We can call this activity evaluation research.

When many people think about evaluating the effect of an ICT innovation, they think of "asking the students", usually by giving them a survey. While students' perceptions about the merit of an ICT innovation are valuable, they are only one source of data, and relying on perceptions alone can give a false impression. For example, in a study where groups of students created their own interactive videos for language learning, half of the students stated that they had not learnt any language through this process. Seen on its own,

this perception may have convinced the teaching staff to discontinue the approach. However, other evidence from video 'out-takes' showed that student teams were indeed engaging deeply with the language, but because of the challenge of learning new technology, they were unaware of this.

Systematic evaluation research needs to go beyond perceptions, and needs to have a clear purpose. A useful distinction is between formative and summative evaluation. Formative evaluation focuses on improvements to products and processes which are being developed, while summative evaluation focuses on the effectiveness of the finished product. Formative evaluation doesn't only concern itself with the ICT product, but also with the learning processes of students and our performance as teachers.

To summatively evaluate the effectiveness of ICT on student learning, we first need ICT which works in the way that it should. We also need to be clear about the type of learning the ICT is designed to achieve. This means we must be aware of research on student learning.

Each individual teacher in a discipline has their own predisposition towards a favoured teaching approach, and their own beliefs about learning. These beliefs often reflect the 'traditional' way that

their discipline is taught paying scant attention to the scholarship of teaching and learning.

For example, in a study of online enhancements to a basic botany course, it was found that the online resources were used heavily, and were found to be valuable by students. However, upon deeper investigation, it was found that the resources reinforced the surface-learning nature of th course, contrary to the intentions of the teaching staff.

You must understand and be comfortable with your personal paradigm of teaching and learning. Within this paradigm, you should be able to articulate why you designed the ICT in the way that you did. It is then much easier to make judgements about how well the ICT performed.

Research paradigmsIn the same way that we need to be aware of our paradigm of teaching and learning, we also need to be aware of our preconceptions about research. As academics, we work within particular research traditions, and these may limit our capacity to evaluate the effectiveness of ICT on student learning. For example, a medical scientist might want to set up an experimental evaluation study, with 'equivalent' treatment and control groups.

Reeves[1] has identified a range of methodological deficiencies in experimental studies, and has suggested that qualitative approaches are more appropriate for the complexity of tertiary student learning supported by ICT.

On the other hand, a social scientist may want to carry out a qualitative study, telling the 'story' of the students in the class. These are also problematic, because they focus on describing what happens, but often without any judgments being made about areas which need change. Purely descriptive studies may be appropriate when we don't understand anything about the phenomenon being studied, but this isn't the case with ICT.

Reeves[2] proposes a pragmatic approach to evaluating the effectiveness of ICT. Instead of comparing 'things' or describing 'things', it is more appropriate to try to discover *how* things work in a particular learning context, using a mixture of qualitative and quantitative sources of data

An apocryphal example relates to an interactive videodisk in the early 1990s. Students using this videodisk to study Physics were found to perform significantly better in exams than students in previous years. However, when a subsequent researcher went to ask the students what happened, they said that the videodisk was so bad that they had to get together in study groups, and go to the library together.

We encourage you to question whether your disciplinary research paradigm is applicable to evaluations of the effectiveness of ICT on student learning. It is preferable to 'step back' from your traditional approach, and, instead of focussing on methodology, to focus on questions to ask, and how best to get answers to these questions. This 'pragmatic' approach is the focus of the rest of this document.

Evaluation and educational design

There should be a close relationship between the educational design of a learning environment and evaluation. Evaluation is an integral part of the **design**, **develop**, **evaluate** cycle of production.

Some evaluation models explicitly map evaluation activities to phases of the development process. As one example, the Learning-centred Evaluation framework[3] has four phases:

Analysis and design: analysing the curriculum, analysing teaching and learning activities; and specifying the behaviour of the innovation.

Development: finding out if the innovation works in the way it was designed, and what is needed to improve it (closely related to formative evaluation).

Implementation: evaluating the effectiveness and viability of the finished product (closely related to summative evaluation).

Institutionalisation: evaluating the effects of ongoing use of the innovation within the institution.

The four phases help you to 'position' your ICT innovation. An evaluation study does not have to examine each phase, but often includes questions related to several phases. In particular, the distinction between formative and summative evaluation often blur.

Advantages and disadvantagesWith this background material in mind, we can start to look at the process of carrying out evaluation research on an ICT innovation. A simple view of the process is to:

work out some questions to ask;

decide who should be asked – the participants;

determine the sources of data to be used;

develop an evaluation matrix;

collect and analyse the data.

The process is discussed in more detail below.

Questions

You need one or two broad evaluation/research questions, which capture the essence of what you want to find out. This relatively difficult task requires clarity of thought about what you want to achieve. However, it is better to think through your questions at the outset, rather than finding out at the end that your questions are inappropriate.

Some people ignore this part of an evaluation. Instead, they feel that they have a methodology (a survey instrument, for example), and they simply apply this methodology to the evaluation context. The danger of this approach is that the evaluation may not consider highly relevant information, and the results may be misleading. Reeves[2] claims this is "akin to a workman claiming to be a "hammer carpenter", stating that the saw and the screwdriver hold little interest".

The best evaluation questions are open-ended, 'how' questions. For example:

How effective are the computer conferencing activities which have been incorporated into this course?

What is the nature of learning processes used by students?

Which factors are important in the design of a learning environment which fosters teamwork?

How can the course be modified to enable students to learn more deeply?

A problem with the open-ended nature of broad evaluation questions is that it may be difficult to answer them. That is why it is helpful to develop a larger number of specific evaluation questions.

These questions should be *answerable* - we should be able to identify the evidence which is needed to answer each question.

Experienced evaluators have acquired the skill to implicitly ask these questions, but novices need some assistance. Evaluation frameworks, such as the Learning-centred Evaluation Framework[3], provide a scaffold to assist novices in developing answerable questions. Some examples of specific evaluation questions are:

How does the approach of the lecturer influence the students' use of the discussion forum?

What is the nature of the teamwork which occurs?

How could the XXX approach be improved?

How do students use the online environment in order to learn?

Participants

The participants in an evaluation study are those who actually provide the data. These are typically students and teaching staff, but may also include technical staff, among others.

Sources of data

There is a range of qualitative and quantitative techniques (methods) which can be used to obtain evidence to answer the specific evaluation questions. Interviews and surveys are the most common.

The LTDI Evaluation Cookbook[13] provides an excellent list of available techniques.

Evaluation matrix

Typically, there is more than one source of evidence which can provide answers to each question. Evidence from a number of sources contributes to the validity of the study. The evaluation matrix draws together the questions, participants and data collection methods. A simple example of an evaluation matrix is shown in Table 1. The matrix helps you to gain an overview of the process, and to plan the rest of the study.

Collect and analyse data

Many people have the research skills to collect and analyse the data generated through the

evaluation matrix, but these skills are beyond the scope of this starter guide. It may be appropriate **to employ suitable staff to carry out this part of the evaluation study.**

- *Practical and simple tips for your evaluation*Take small steps: don't try to understand everything at once.

- Be a reflective practitioner: remember you are trying to improve your students' learning experiences and, indirectly, your teaching.

- Use a cycle of understanding and improving: think formatively rather than summatively.

- Question the assumptions underpinning your personal paradigm of teaching and learning.

- Question the applicability of your disciplinary research paradigm to evaluations of the effectiveness of ICT on student learning.

- Document the design of your ICT, explaining why this is likely to lead to the learning outcomes you require.

- Think of ways in which your evaluation can go beyond student perceptions.

- Use experimental techniques cautiously, as it is difficult, if not impossible, to keep elements of the learning experience constant between groups and different learning contexts.

- Evaluate student learning in the whole teaching context, not just the ICT itself.

- Examine not only what students learn, but how they learn: reflect on the relation between learning process and outcomes.

- Focus on questions to ask, and how best to get answers to these questions.

Use an evaluation matrix to organise and manage the evaluation study.

Conclusion

Learning-centred evaluation is difficult to do and takes lots of time. Learning is a complex human activity – learning with ICT is even more complex and uncertain. However, investigating the relationship between what we learn and how we learn with ICT should take us a long way to understanding this complex computer-mediated activity.

We need to treat teaching and learning in much the same way as we treat research, where we seek to understand what is happening, and critically reflect on what we have learnt.

References and Resources[1.] Reeves, T.C., Pseudoscience in computer-based instruction: the case of learner control research, Journal of Computer-based Instruction 20 (2), 39-46, 1993.

[2.] Reeves, T.C., Established and emerging evaluation paradigms for instructional design, in Instructional Development Paradigms, Dills, C.R., and Romiszowski, A.J. Educational Technology Publications, Englewood Cliffs, New Jersey, 1997, pp. 163-178.

[3.] Phillips, R., Bain, J., McNaught, C., Rice, M. and Tripp, D., Handbook for Learning-centred Evaluation of Computer-facilitated Learning Projects in Higher Education, Committee for University Teaching and Staff Development Project, 2000. http://cleo.murdoch.edu.au/ projects/cutsd99/handbook/handbook.htm

[13.] Harvey, J., Evaluation Cookbook, Learning Technology Dissemination Initiative, 1998, http://www.icbl.hw.ac.uk/ltdi.

Acknowledgements