The prospects for summative evaluation of CAL in HE

Published: ALT-J (Association of learning technology journal) (1997) vol.5, no.1 pp.33-39

Stephen W. Draper
Department of Psychology
University of Glasgow
Glasgow G12 8QQ U.K.
email: steve@psy.gla.ac.uk
WWW URL: http://www.psy.gla.ac.uk/~steve

Abstract

Many developers and evaluators feel an external demand on them for summative evaluation of courseware. They feel they are being asked to prove that the software "works", to show that it is cost effective, that it is durable, that it is worth the price to the purchaser. However as soon as you start to attempt, or even seriously to plan, such studies the problems with this aim emerge. One is that the CAL may not be used at all by students if it is not made compulsory. If you measure learning gains, how do you know you are measuring the effect of the CAL or of the motivation in that situation?

Such issues are the symptoms of the basic theoretical problem with summative evaluation, which is that CAL does not cause learning like turning on a tap, any more than a book does. Instead it is one rather small factor in a complex situation. It is of course possible to do highly controlled experiments: e.g. to motivate the subjects in a standardised way. This should lead to measurements that are repeatable by other similar experiments. However they will be measurements that have little power to predict the outcome when the CAL is used in real courses. Hence the simple view of summative evaluation must be abandoned.

However it is possible to gather useful information by studying how a piece of CAL is used in a real course and what the outcomes were. Although this does not guarantee the same outcomes for another purchaser, it is obviously useful to them to know that the CAL has been used successfully one or more times, and how it was used on those occasions. Such studies can also, as we have demonstrated, serve a different "integrative" rather than summative function by pointing out failings of the CAL and suggesting how to remedy them.

Introduction

Summative evaluation is evaluation done after software design and production is complete in order to establish its performance and properties. A prototypical case would be the tables produced in the consumer magazine "Which?" comparing a considerable range of properties of alternative available machines (e.g. washing machines) that they have measured in their own trials. Thus summative evaluation is not only done after production, it is typically about comparative measurements done to assist decisions about purchase.

Many developers and evaluators feel an external demand on them for summative evaluation of courseware. They feel they are being asked to prove that the software "works", to show that it is cost effective, that it is durable, that it is worth the price to the purchaser. This is seen as a matter of testing the software, as in most software projects. Thus tests are done by using the software, and measuring various outcomes of its use (e.g. how people think about it, what is learned). Sometimes the software's performance is compared with some alternative e.g. no software, traditional teaching, etc.

However we know much less about the ingredients of successful teaching delivery than we do about washing clothes, and this has important consequences for what we can learn from measurements and what we in fact want to find out.

Symptoms of problems with the obvious approach

In our work on the TILT project (Doughty et al. 1995), we were soon struck by features that cast doubt on the sense of our doing evaluation of this kind (Draper et al. 1994).

1. One is that the CAL may not be used at all by students if it is not made compulsory. (Here I shall use the term "CAL" to refer indiscriminately to any comptuer software that might be introduced by teachers to support learning.) This draws attention to the crucial role of motivation. If you measure learning gains, how do you know you are measuring the effect of the CAL or of the motivation in that situation? Certainly in the case mentioned, the CAL alone produced no learning because it produced no usage: motivation created by a teacher was crucial.
2. Another issue is that of the actions of teachers, for instance engaging students in a socratic dialogue based around the CAL. Obviously any learning gains would be affected, and probably dominated, by the teacher's skill. But to evaluate software in the absence of teachers is to measure a different situation than the most common one in HE, and furthermore one that would not get the most from the software: hence it is neither realistic (valid) nor constructive.
3. Unstable student study strategies: we have seen student study strategies such as note taking radically changed by short remarks by the teacher, far more so than in a lecture. At least at the present time, it seems that CAL does not usually elicit a stable study strategy while teachers can and do influence it in a big way. Since study strategies have a large effect on learning outcomes, again it seems beside the point to look for measurements independent of these: rather the point would be to discover which study strategy is best for each piece of courseware and how to ensure that students adopt it.
4. Although some software is designed to be used once and never referred to again, a lot of courseware is like textbooks and is intended to serve as reference and revision material as well as or instead of as primary exposition. That means that the relevant tests of learning must be delayed until after the exam. However it is then hard to tell how much, if at all, the students depended on the courseware as opposed to alternative resources such as books. Any evaluation will be telling you not about the properties of the software but about how the overall set of resources and student activities performed.
5. As a corollary to this, if in fact a student finds the courseware useless they are likely to compensate by relying more on alternative resources. (We might call this self-monitoring and correction "auto-compensation".) In universities, poor teaching may often be masked by this, and final performance relatively little affected. On this view, studying the effect of courseware in isolation is unrealistic, but overall performance depends mainly on the students' self-management rather than on any one resource. One could however attempt to study which resources students use and value (Brown et al. 1996).
6. Halo effects can also be important, where a teacher's attitude to the technology may strongly affect students in either positive or negative directions. While we have seen marked effects of this kind on student attitudes, whether this matters for learning depends also on whether student attitudes affect their use of the CAL: if they have no alternative resource, it may not matter.
7. Similarly Hawthorne effects may occur, where the act of doing the evaluation may affect students by making them feel more valued, pay more attention to the CAL and perhaps the subject matter it deals with. In addition to an effect on the learners' attitude, pre-tests given as part of an evaluation may well improve learning by communicating to the students what they should try to learn from the material. Furthermore priming students to activate the relevant part of their theoretical knowledge is known to have a big effect in improving learning outcomes from lab classes and simulations, whether or not this is done by the evaluation or by a deliberate part of the teaching (e.g. pre-lab exercises). All you can evaluate is the combined effect of the CAL, the evaluation, and the whole of the associated teaching and learning resources. This is fine from the point of view of improving learning and teaching, and indeed evaluation should probably be a permanent part of practice, but it again undermines the view that the effects of CAL can be studied as an independent topic.

CAL is only part of an ensemble

The fundamental point is that CAL does not cause learning, and in fact is not a major cause at all. Learning results from the combined effect of many important factors and typically, in universities, from multiple resources. Any realistic study or evaluation measures the combined effect of an ensemble. This does not mean evaluation studies are impossible, but it does mean we need to think out what we really want to discover. We cannot expect to treat CAL like a washing machine: as a simple device whose performance can be measured once in a standardised situation which will then tell us all we need to know to decide whether and how to use it.

If we remember that testing a piece of CAL is essentially the same as testing a textbook would be then this seems obvious. It is also like considering the question "is the 9:30 Glasgow-Edinburgh train good for getting to Edinburgh?" It is possible to imagine that there could be something uniquely good or bad about that train and not others, but in fact usually the important factors are not the details of the train itself but how it fits into people's overall travel needs and plans. People only use trains as part of wider plans, and trains are mainly good or bad to the extent that they fit (or don't) into the success of these wider plans. To do meaningful evaluation of CAL, we have to understand learners' wider plans and study them: what they are, what the main factors are that influence their success, where CAL fits into this.

The issues listed above are the symptoms of the basic theoretical problem with summative evaluation, which is that CAL does not cause learning like turning on a tap, any more than a book does. Instead it is one rather small factor in a complex situation.

Integrative evaluation: the actual utility of "summative" evaluation

In the TILT project, we performed many evaluations on completed software in the classroom. We found that our evaluation reports were often useful to teachers, but not for summing up the properties of the software so much as for identifying specific problems in the case being studied that the teachers would use to make changes, usually not to the software, but to some other aspect of the delivery e.g. modifying how they introduced the software, adding a section to a lecture. We thus realised that the value of our evaluation was not as summative evaluation of the software, but as formative evaluation of the overall teaching and learning situation. We called this "integrative evaluation" (Draper et al. 1996) because most of the changes made concerned improving the embedding or integration of the software into the rest of the surrounding delivery.

Experience elsewhere seems to bear this out: often when CAL is first introduced there are some problems, particularly if students feel they have been set adrift without the right kind of support and guidance. If evaluation is taken as measuring how good or bad the software is, it would have to record a partial failure and would have to refrain from making contributions to improvements. However this is not how sensible participants actually use it: they use the evaluation to alert them to problems and quickly introduce improvements. The next time the course is given, the evaluation can often verify the improvement, but its most important role has been identifying problems and necessary improvements. Conversely when participants treat an evaluation as strictly summative, problems can result. In an unpublished case, an evaluation was commissioned to help the institution decide whether to adopt a substantial CAL package as part of a course. The report was then interpreted strictly as supporting the adoption, but all its advice on the crucial integration issues it identified was overlooked and no further evaluation was permitted. Standard student feedback later indicated substantial quality problems on the resulting course, apparently around those integration issues identified in the original study.

What we really want, and what we can do

Evaluation is only worth doing if it serves some purpose and leads to some useful action. The previous section described how often evaluations on completed software are used to identify integration issues and so improve the total teaching and learning: and we might call this "integrative evaluation". Are there, then, any other goals that a summative evaluation might be required to help us with?

Yes: the goals informally expressed as "Are we going to use it? And how are we going to use it?". An important decision that should be supported by information from evaluations, in education as in consumer purchases, is whether or not to buy and adopt some product, in this case a piece of CAL. Relevant to this would be answers to whether students do learn in courses adopting the CAL, whether this worked in other institutions, and particularly in my own, what it costs in resources to run the course with the CAL, and what you need (and need to know) in order to run such a course successfully i.e. a description of the whole teaching delivery including auxiliary materials would be very desirable.

The fact that learning outcomes depend on the combination of many factors of which the CAL is only one means that no single study can prove that the CAL will work in any other situation e.g. it might work for the authors, but not work in the different context of a new adopter. However, while this does mean that certainty is beyond reach, it does not mean that evaluation is worthless for this. Imagine what you would find useful and persuasive as evidence about whether to adopt a piece of CAL. The first important thing to discover is whether it has ever been used successfully i.e. with satisfactory learning outcomes. Even in fields much better understood theoretically,

such as building aircraft, the first use is crucial to demonstrate that no crucial mistake has been made: you do not expect to be the first person on an aircraft never used before. Unlike aircraft, however, the performance of CAL depends a lot on the surrounding context of use, so tests with real students as part of a real course are much more convincing than tests in a lab with paid subjects. Thus even though certainty may be beyond reach, such tests that the CAL can be successful are an important reason for and outcome of summative evaluation. Beyond that, the questions (see above) shade into issues of how best to use it. The more that these issues are identified explicitly and made available in reports or auxiliary material for teachers, the better for those deciding whether to adopt it. Thus pure demonstrations of possible success and outcomes from integrative evaluation can together serve the essential underlying goal of supporting decisions about whether to adopt the software.

Experiments

A related issue is the use of controlled experiments. Summative evaluations on consumer goods are based on controlled experiments e.g. using the same standard load of dirty clothes and the same detergent for each washing machine compared. On the other hand, the useful results of integrative evaluation are often (though not always) the result of open-ended observation or student feedback, identifying factors that had not been foreseen and so not systematically measured. Furthermore, all the points above about the factors likely to be important in affecting outcomes suggest that we do not know enough to control all the relevant factors. Few if any experiments for instance attempt to control the halo effect or student's uncertainty about study strategies with CAL. Can experiments be used meaningfully at all in CAL evaluation?

This remains arguable. Clark (1983) has argued that no meaningful experiments on whether learning is affected by the medium of instruction have been or could be done, because other factors more plausible as the causative agent vary with the medium. While hotly debated in the literature (Ross; 1994), his arguments have not been conclusively rebutted and they apply also to the use of experiments in evaluation. They would apply most strongly to prevent us from drawing generalisations such as that a piece of CAL will always be successful. Nevertheless, some experiments seem rather convincing e.g. MacDonald & Shields (1996), particularly when alternative ways of teaching are directly compared e.g. lectures, CAL with and without special worksheets. Note that such experiments can in part serve an integrative role by yielding information on how best to use the software: there is no exclusive association between evaluation methods (experiment vs. open-ended observation) and the evaluation goal (summative or integrative). In the end, experiments are probably like other studies of CAL: they can show that the CAL was definitely part of successful learning in one case, and if no other cases have been reported then this is in favour of the CAL. On the other hand, we remain aware that many factors may have been important, and the report never describes them all. It could always be that one of these factors was crucial, but will not be present if you try to use the CAL in your own teaching e.g. perhaps freeing up lecturers meant that while in the lab supervising CAL use, they performed tutorial interactions with students that were crucial, just to have something to do.

Conclusion

We can do evaluation that is summative in some senses but not in others. We can do evaluation at the end of the design cycle on completed software that will not be further modified. We can do evaluation that provides evidence relevant both to deciding whether to use that piece of courseware for teaching, and also on how best to use it. This latter evidence might be from careful comparative experiments or more formative style work that detects unforeseen problems that turn out to be important in successful use of the software. But there is no prospect of doing evaluation that sums up a product once and for all, measuring its essential properties in a way that will represent and predict its performance in all other contexts.

Acknowledgements

This paper stems from work on the TILT (Teaching with Independent Learning Technologies) project, funded through the TLTP programme (Teaching and Learning Technology Programme) by the UK university funding bodies (DENI, HEFCE, HEFCW, SHEFC) and by the University of Glasgow. The ideas come from collaboration with other members of the evaluation group, paritcularly Margaret Brown and Erica McAteer. The studies mentioned here could not have been done without, in addition, the active participation of many members of the university teaching staff to whom I am grateful.

References

Brown, M.I., Doughty,G.F., Draper, S.W., Henderson, F.P. and McAteer, E. (1996) "Measuring Learning Resource Use." Computers and Education vol.27, pp. 103-113.

Brown, M.I., Draper, S.W., Henderson, F.P. and McAteer, E. "Integrative evaluation as an agent for change." ALT-J (this issue).

Clark, R.E. (1983). "Reconsidering research on learning from media" Review of Educational Research, vol.53 no.4 pp.445-459.

Doughty,G., Arnold,S., Barr,N., Brown,M.I., Creanor,L., Donnelly,P.J., Draper,S.W., Duffy,C., Durndell,H., Harrison,M., Henderson,F.P., Jessop,A., McAteer,E., Milner,M., Neil,D.M., Pflicke,T., Pollock,M., Primrose,C., Richard,S., Sclater,N., Shaw,R., Tickner,S., Turner,I., van der Zwan,R. & Watt,H.D. (1995) Using learning technologies: interim conclusions from the TILT project TILT project report no.3, Robert Clark Centre, University of Glasgow ISBN 085261 473 X

Draper,S.W., Brown,M.I., Edgerton,E., Henderson,F.P., McAteer,E., Smith,E.D., & Watt,H.D. (1994) Observing and measuring the performance of educational technology TILT project report no.1, Robert Clark Centre, University of Glasgow

Draper,S.W., Brown, M.I., Henderson,F.P. & McAteer,E. (1996) "Integrative evaluation: an emerging role for classroom studies of CAL" Computers and Education vol.26 no.1-3, pp.17-32 and Computer assisted learning: selected contributions from the CAL 95 symposium Kibby,M.R. & Hartley,J.R. (eds.) (Pergamon: Oxford) pp.17-32

MacDonald,Z. & Shields,M. (1996) "Effective computer-based learning of introductory economics: some results for an evaluation of the WinEcon package" submitted to Journal of Economic Education

Ross,S.M. (1994) "Delivery trucks or groceries? More food for thought on whether media (will,may, can't) influence learning: introduction to special issue" Educational Technology Research and Development vol.42 no.2 pp.5-6