

You have carried out your evaluation study and got your results – what do you do next? Has the evaluation fulfilled your original aims? Have you found out what you wanted to find out? Do you need to follow-up some of your findings?

Everything was wonderful?

It is sometimes tempting to be selective when carrying out your own evaluation because you have invested so much effort in setting up an implementation and you feel sure that it must have worked well. Watching students working enthusiastically through a package, obtaining positive feedback from a questionnaire or finding evidence to support a gut feeling that an intervention was successful is always a positive if not a reassuring experience.

It is worth bearing in mind that there are always some aspects of an intervention that can still be improved and that sometimes one teaching method can work well with one group of students and not quite so well with another group. Asking students for suggestions for ways in which the intervention could have worked better for them can elicit a lot of surprising and constructive comments. Following these comments up will also demonstrate to students that you value their opinions as well as giving them more of a feeling of responsibility towards their own learning.

If you feel that the data from your evaluation study was biased a bit too much towards the positive side, you might want to consider the way in which your study was conducted. For example, did you only ask questions that would elicit a positive answer? Do you think that the feedback was a product of the student sample selected and might another group of students from that class have produced very different feelings? Were your students just trying to please you? This is not to say that you should question the validity of all your positive data, but that if you are wanting to try and improve an intervention, then you should watch that you don't adopt too much of a 'feel good' strategy.

What if it didn't go so well?

At the other extreme, it can be very disheartening to obtain a stream of negative comments from students. This is an added problem when you are evaluating someone else's pet project. It is surprising how much something as fundamental as having very limited access to machines can colour students' general feelings about using computers and can consequently produce a lot of negativity about anything that is computer based. Students' feelings can also be influenced by the timing of the evaluation study, their previous learning experiences and even how well they related to the person carrying out the evaluation study.

Perhaps your questions have been phrased in such a way as to elicit negative comments or have focused on the negative aspects of an intervention. Asking students about the most useful or their favourite parts of the software can encourage students to think about the positive aspects of a piece of software instead of dwelling on the negative areas.

A well planned evaluation study should aim to focus on a few key areas of concern within an implementation. These might relate to the software content and structure, the practicalities of using the software in your institution or the way in which the materials might fit into a specific course syllabus.

What about the unexpected results?

Sometimes an evaluation study can produce not just positive or negative results but completely unexpected findings. This can happen when you have adopted a less structured approach to your evaluation study. Perhaps you find that your students are working in groups instead of individually or spending an excessive number of hours going through a set of self-assessment questions until they get 100% or if you haven't given them a specified task, only flicking through the parts which interest them. You may want to think about whether or not these outcomes will add or detract from your original aims and objectives of the intervention.

What if you want to make changes to the software?

If problems are identified within the software then these can be reported back to the developers. Did your students have difficulties navigating through the package? Was the layout or structure confusing? - in what way? Did they have difficulty in using particular parts of the package? Were graphics/text/colours used appropriately? Was the subject content factually correct?

Even if you aren't carrying out a software product trial, developers will generally welcome constructive comments relating to their products and might be prepared to make changes to the content based on recommendations made by a number of academic staff. However, identifying limitations does not necessarily mean that changes have to be made to the software itself, you may want to consider how some sections from other packages or text based materials might be used to provide additional support for students.

Carrying out an evaluation with a group of students prior to a full implementation is useful not only in identifying problems relating to the software but also to the practicalities of using the software within your particular institution.

Hardware problems?

Hardware problems might be a result of how and where a piece of software is being used. Even trying to organise an evaluation study can demonstrate the practicalities of using technology within your institution. How easy is it to book ten machines in your department for students to use? Is it a problem for students to obtain access to a machine outwith time-tabled hours? How easy is it for students to load up and use the relevant software? Can a full class of students access one particular piece of software at one time?

It is advisable to try and alleviate any procedural problems prior to the start of the course in which the technology is embedded rather than leaving things to the last minute and hoping for the best. For example, providing students with a list of instructions on how to book a machine, how to access and make use of the relevant software can help to provide support and encouragement to your students. If there is a problem with a whole class trying to download a piece of software at one time, it might be appropriate to try and stagger the times at which the students log in to the system and this might involve rescheduling sessions well in advance.

Even if all of your students are able to access machines at one time, you might encounter practical problems in using the software. For example, was there a problem loading any of the sections? Were the appropriate plug-ins available? Could the students obtain a print out when required? Identifying this type of problem does not necessarily mean that you can't make use of the software. A number of alternatives could be suggested to your students: you might wish to warn students if certain graphics or video clips are likely to prove slow to load on old machines, and you could suggest that they miss that particular section out. If it is important that students see what is contained in a video clip then you might want to show the clip on a faster machine in a tutorial. This would also give you an opportunity to discuss any other issues that might be raised during their time using the package.

Problems with the subject content of a package

Your evaluation study might have identified that students have problems with some parts of the subject content of a package. Assuming that the problem is not neither the accuracy of the material or the layout of package, it is advisable to identify exactly where the problem lies. This might mean following up an initial evaluation study with a more focused investigation. Which students were having difficulty – the top/bottom or all of the class? Was this difficulty with one part of the software only? Did the students not have the pre-requisite knowledge to make best use of the package? Sometimes it can be sufficient to tell students that there are online help facilities such as a glossary available within a package. If there isn't one available then you could consider producing your own help sheets which are relevant to the tasks you have set for your students. If only some were experiencing problems, asking students to work in groups while using the software allows them to help each other.

Should I bother using this software?

If a large proportion of students in your evaluation study reported problems with a piece of software, then you might want to think about the usefulness of including the package in a course at all. But if you feel that the package includes material that is very difficult or costly to explain in any other way then perhaps you could review how the material is embedded into the course. In the same way as you might refer to certain chapters in a textbook you might want to encourage students to use only certain parts of a software package. Providing an introductory tutorial prior to the students' use of the software can set the scene for the subject material in a CAL package, as well as allow you to fill in any gaps in the software subject content.

The students didn't seem to want to use the software

You might have found that fewer students than you had expected have made use of software particularly if materials were made available within an open access centre. This could have been because access to the materials was limited or problematic or that the software was just viewed as being an optional extra. Any students that provide feedback as part of this kind of study, should be considered to be a self-selecting sample. As a result, this data is likely to be different to results obtained from a more representative sample. Although you might have obtained very positive feedback from these students, you should still take this lack of uptake by the majority of the class into account. If students feel that a new piece of technology doesn't warrant either your enthusiasm or the full integration into a course then they are going to be less likely to spend time making use of the materials.

What about the rest of the course?

An intervention is only one small part of a course. If you have carried out an evaluation study of the software in isolation, then this does not become such an issue until you are thinking about the way in which the materials are going to be used during the next year. The way in which the material is embedded into the course will influence the overall effectiveness of a piece of software in achieving specific learning objectives. For example, in a preliminary evaluation study, you might have found that a simulation package was very effective in encouraging students to develop a problem solving approach but used in a subsequent class where the end of module assessments encouraged factual recall, students adopted very different learning strategies.

If you have carried out an evaluation study of part of an ongoing course, it is worth bearing in mind the impact that this might have on your students' perception of the focus of your attention. Do they start working harder for this part of the course to the detriment of the remaining parts of the course? Does this part of the syllabus warrant such attention? Is the improvement in the end of module marks a result of the intervention or the level of importance attributed by the lecturer carrying out the evaluation? Can you expect the same results next year when perhaps you and/or the lecturer are focusing on a different part of a course?

Reflections on your evaluation

Evaluation is part of an ongoing process and not just a one off event. The findings from an evaluation study should be used to make improvements to the teaching and learning within a course. If students make any recommendations, try to take these into account and to report the resultant changes back to them.

The process of planning and carrying out an evaluation study is a useful exercise in that you start to focus on specific learning objectives. This in turn, gives you an opportunity to reflect on the most effective ways in which your students might attain these objectives.

Jen Harvey

Implementation Support Consultant,
LTDI, Heriot-Watt University.