

EVALUATION COOKBOOK

Editor Jen Harvey

The Learning Technology Dissemination Initiative is funded
by the Scottish Higher Education Funding Council

First Published 1998 by

Learning Technology Dissemination Initiative
Institute for Computer Based Learning
Heriot-Watt University
Edinburgh
EH14 4AS

Tel: 0131 451 3280
Fax: 0131 451 3283
e-mail : ltidi@icbl.hw.ac.uk
URL: <http://www.icbl.hw.ac.uk/ltidi>

ISBN 0 9528731 6 8

If you would like further copies of the LTDI Evaluation Cookbook please contact LTDI.

Original Authors (c)1998 LTDI

All rights reserved. Apart from fair dealing for the purposes of research or private study, or criticism or review, as permitted by the Copyright, Designs and Patent Act 1988, this publication or any part thereof may not be reproduced stored or transmitted in any form or by any means without the prior written consent of the publisher, LTDI. Except that (a) for courses of instruction or distribution within UK higher education institutions this publication or parts thereof may be copied without charge, provided that the source and LTDI's copyright are indicated in any such reproduction and (b) authors retain their personal right to re-use the material in future versions or collections of their own work.

Whilst every effort has been made to ensure the accuracy of the information in this book, LTDI wishes to emphasise that the LTDI cannot accept any liability for any errors which remain. Further the views expressed in each of the chapters are primarily those of the stated authors and whilst LTDI endorses the value of all the contributions, the views expressed are not necessarily those of LTDI. Please notify us of any factual errors, for correction in future editions. We would also welcome any other suggestions, additions or amendments that would enhance this publication.

Background

The original idea for the cookbook came from the LTDI consultative group. Their suggestion was that we should try and produce a practical guide that would be aimed at lecturers who were interested in evaluating materials for their effectiveness in achieving specific learning objectives.

From this initial concept, we hope that we have now developed a resource from which readers can pick and mix from a range of ideas and suggestions and through this process, design an evaluation study to suit their own specific needs.

The development of the cookbook

Cookbook contributors were selected on the basis of their wide experience in carrying out evaluations of teaching and learning interventions and we feel that this is reflected in both the content and the approach adopted within each of their cookbook sections.

In order to encourage collaboration between the authors, who were based in universities as far apart as Aberdeen and Auckland, the LTDI set up a private WWW site and mailbase discussion group. Once contributors had completed their section(s), each of these was then made available for review from the WWW page. Any comments were then fed back either directly to the author, to LTDI or to the whole group. Authors were encouraged to feedback comments on all the sections.

In addition to this, it was decided to allocate each of the contributors to a group of three and to ask them to make a more detailed review of the other two author's work.

We are making the WWW site used during the development of the cookbook available to you and would welcome any comments, feedback or ideas you would like to make relating to the 'cookbook project'.

The WWW site is available at
<http://www.icbl.hw.ac.uk/ltdi/cookbook/>

Preparation pages

Introductory sections have been included to provide a framework to the planning and preparation process involved prior to carrying out an evaluation. These aim to encourage you to think, in more detail, about who the evaluation is for, what you are going to be evaluating and how best you might carry out such an evaluation study.

Recipe pages

Each recipe comprises a summary of the main uses of that particular method, a step by step guide to the time, resources and process likely to be involved, as well as a set of cook's hints relating to each stage of the process. Links to other relevant pages in the cookbook are also included.

Information pages

The recipes are interspersed with information pages that aim to provide some basic practical suggestions and advice, applicable to a range of different evaluation methods. Links are made from recipe pages to any relevant information pages.

Tasting, Refining and Presentation pages

The final sections of the cookbook encourage you to think of your evaluation study as an ongoing process used to make improvements in teaching and learning. Guidance is provided to encourage you to reflect on ways in which you can act on your results and/or write up your findings in an evaluation report.

Serving Suggestions

The serving suggestions sections are included to demonstrate some of the cookbook evaluation methods put into practice. These short exemplars outline the aims and objectives of various evaluation studies, the main findings from these studies and some reflections on these findings.

Preparation

Why Evaluate?	8
Directing your evaluation?	9
Selecting a methodology?	12

Recipes

Quick guide	15
Checklists	16
Concept maps	22
Confidence logs	26
Cost - effectiveness	28
Designing experiments	32
Ethnography	34
Focus groups	38
Interviews	40
Nominal group technique	44
Pre and post testing	48
Questionnaires	50
Resource questionnaires	54
Split screen video	56
Supplemental observation	60
System log data	62
Trials	64
LTDI chocolate cake recipe	66

Tasting, Refining and Presentation

After the evaluation	68
Evaluation report outline	71

Serving suggestions

Checklists	73
Concept maps	
Confidence logs and questionnaires	
Ethnography including observational techniques	
Focus groups, user logs and questionnaires	
Nominal group technique	
Pre and post testing and tracking log data	
Questionnaires and interviews	
Trials	

Information Pages

Isolation or Integration	18
So you want to use a likert scale?	21
Taking learning styles into account	24
Statistics questions	30
Selecting your student sample	31
Working with groups of staff and students	36
Practical hints when interviewing	42
Pre and post testing?	46
Questionnaires - some advantages and disadvantages	52
Some guidelines for writing good questions	53
Transcribing evaluation data	58
A pedagogical toolkit	63

References

83

List of contributor's addresses

87

Guest Cooks

Gayle Calverley, The University of Hull
Grainne Conole, University of North London
Philip Crompton, University of Stirling
Kirsty Davidson, Napier University
Judith George, Open University
Judy Goldfinch, Napier University
Cathy Gunn, University of Auckland
Chris Jones, Liverpool John Moores University
Liz Leonard, University of Glasgow
Ann Lewins, University of Surrey
Gaye Manwaring, Northern College
Patrick McAndrew, Heriot-Watt University
Erica McAteer, University of Glasgow
Hamish Macleod, University of Edinburgh
Colin Milligan, Heriot-Watt University
John Milne, Aberdeen University
Nora Moge, Heriot-Watt University
Robert Newton, The Robert Gordon University
Martin Oliver, University of North London
Robin Shaw, University of Glasgow
Helyn Thornbury, University of Strathclyde
Gwen van der Velden, University of Kent

This publication would not have been possible without the following people:

The LTDI consultative group who suggested we produce this publication.

SHEFC for funding the project.

The cookbook authors who contributed their time, ideas and experience at all stages of the project.

Nora Moge and all the other people based at the Institute for Computer Based Learning who gave such useful and constructive comments.

Phil Barker for all his help and patience in setting up the WWW site and editing the cookbook.

Mary Cuttle for the LTDI chocolate cake recipe.

Graphic Services, Heriot-Watt University for the design and production.

The ScENet group, the recently restructured Scottish Evaluation Network, for their help and support throughout the project.

“When the cook tastes the soup, it is formative evaluation;
when the dinner guest tastes the soup, it is summative evaluation.”

The Costs of Evaluating

Evaluations are costly. Even the simplest takes precious time from other activities. Apart from the too familiar questionnaire that we usually ignore at the end of some workshop or presentation, a lot of careful thought has to go into constructing a worthwhile evaluation. Then there's teaching time which has to be sacrificed to getting the students to complete questionnaires. And the leisure time you have to devote to interpreting the results and writing them up in a form that others can understand. So why evaluate? Well sometimes you'll be forced to because future funding or support depends on it. But more often you'll evaluate to learn. If there's to be no action taken as a result of the evaluation then, unless you just want material for a paper, the effort is not worth pursuing.

Who Gains?

The starting point for any evaluation is to identify the stakeholders. In some cases the stake is hovering above your project and you're looking for evidence to prevent it being driven home. But usually the evaluation is being conducted to bring some benefit to one of the groups of stakeholders.

Let's start with students, since often they are an afterthought. What are the concerns of the group you're targeting? There are some obvious areas which interest them, from gaining a better education, through issues of accessibility, to the passing of the coming exam. As with all the stakeholders, don't plunge into constructing the evaluation without talking to them and exploring their concerns around the educational intervention on which you're focusing. Then the resulting evaluation will be centred on discovering how the intervention can be improved to satisfy the real aims of the target audience rather than what you decided they should be.

Evaluating for developers is more straightforward. Given that the content is appropriate, the developer is interested in how easy or difficult the user found it to access the material. Were there any bugs? Was the navigation instinctive? Was the text in a suitable font and was it presented in appropriate volume? Was the feedback provided at the right place and did it satisfy the user? And so on.

Lecturers want to know about learning gains and efficiency. Was this a better way of presenting the material than the tutorial or the lecture? Did it free up time for more effective

contact with the student or to conduct research? Are there additions such as handouts which need to be considered to improve the effectiveness of the intervention?

Management need evidence that the time spent on development has led to greater efficiency while maintaining at least the same quality. Does the product justify the costs? Have the students welcomed the change and will the course continue to attract recruits? Have the exam results remained acceptable? Will it help with the TQA? Can fewer resources be devoted to the course than before?

There are usually other stakeholders who have a wider interest in the results of evaluations, especially of computer assisted learning. The Funding Councils, for example, wish to consider whether money is well spent in this area, and though a large external evaluation will usually be conducted to provide the answer, the sum of small local evaluations feed into the decision.

Will it be Worth it?

So, before you embark on an evaluation, ask yourself "why bother?". Who is this for, what is it they want to find out, and what changes will be made when the results are gathered? If the answer to the question "why evaluate?" is that the results will lead to action to improve the teaching and learning within the course or the institution, then all the effort will be worthwhile.

Robin Shaw
TLTSN Consultant,
University of Glasgow

Evaluation studies are fundamentally about asking questions, and then designing ways to try to find useful answers. Studies may concern materials, projects, courses, methods, packages, or systems; in fact, anything that can be asked about in a detailed, structured fashion.

In formative evaluation, information can be transferred back into the original work to both strengthen and move it forward. It is an ongoing, fluid process, used to gauge overall progress and areas needing some attention or change, helping to mould the final article. In summative evaluation, the information is intended to give an overall picture at the end of a stage, often measured against fixed criteria. Summative evaluation provides a fixed point of reference, and it may provide a measure of success or otherwise against original objectives and planned outcomes or it may include reactions from participants to a goal free investigation.

It is crucial to take time at the very beginning to determine which are the “right” questions. Inappropriate or unrealistic questions will lead to unusable or irrelevant data, rather like setting up a computer to perform a complex calculation only to find it was given the wrong formula to start with. But it may also become apparent during a study that some questions are unhelpful and need to be changed, and others added, so build in enough flexibility and open-endedness.

Think about the framework of the proposed study, and how this fits in with the work it is intended to evaluate. The following headings are offered as a starting point, and include suggestions to help determine what aspects are most important in your particular study. The items are given in no special order, but are intended to provoke thought.

What will your evaluation do?

When you plan a meal you know if you want a sumptuous banquet or a snack to eat in front of the TV. You also consider your guests and their tastes as well as the budget and time you have available. Similarly, when you plan an evaluation you must consider the purposes, the interests of those involved and the practical limitations.

Are you:

- ◆ putting a trial software product in front of potential users?
- ◆ doing a preliminary survey to determine a need for a particular service or product?
- ◆ carrying out an information poll for a third party?
- ◆ testing a final system under its real everyday circumstances?

Are you looking at developing a comprehensive, multi-stage evaluation, requiring several smaller self-contained studies? Is there a need for several of these studies at different stages in development or will a single one do?

Who is the evaluation for?

There will probably be several interested parties e.g.

- ◆ those with a direct investment (the stakeholders);
- ◆ those who may be carrying out similar work in the future;
- ◆ those you may want to educate through your work.

In a new course, the key stakeholders may have different concerns.

- ◆ The students may be more interested in a formative evaluation that can address any problems before the end of their course;
- ◆ A lecturer trying out a new piece of software may want to evaluate its potential for transfer to other courses;
- ◆ Senior managers may be interested in comparisons between different courses in terms of completion rates and customer satisfaction;
- ◆ Employers may be interested in the demonstrable skills of those taking the course.

You may not be able to satisfy all the needs but you can try to explain what you see as the main purpose of the evaluation.

Remember too that YOU are making the investment in

performing the study. What type of information is most important for you to meet your goals and objectives? What information will help you to convince key groups of the value of your work? What areas of your work would you like to examine more closely?

Performing an evaluation study is a good opportunity to be able to stand back from the work you are doing and appraise it. Critically there may be specific questions such as “Are my students struggling with the new module design?”, “Are we being cost effective?” or “Are there any specific gaps in the system we haven’t noticed?” a well-designed study can draw all of these concerns together to provide an overall picture.

Can you deal with the practicalities?

What is the size and scale of your evaluation in terms of numbers involved and the timescale? If you have a large number of students you may want to sample their performance and views. If you are evaluating in a number of different contexts you may want to choose varying environments. You may need a quick answer to let you make a decision next week or you may want to analyse long term effects over time.

You will need to consider who will carry out the evaluation. An internal evaluator will understand the nuances of the context but an external person may be more objective. Can you get help? For example, you may be able to employ a research assistant for a few hours to do some interviews or a computer analysis of results. Estimate the time needed for each stage – planning, designing instruments, collecting data, analysing information, making decisions and reporting findings. Make sure you choose the best time to carry out the evaluation – when enough has happened, but not when the respondents are busy with exams. Also consider the timing of your study. Does it have to fit into an external schedule? For example, if you are working with a development team, what is their release calendar? If you are working with students, when is the course delivered? Is the release schedule compatible with the course schedule and is either negotiable? Co-ordinate the focus of the study with the state of the work at the time it is actually going to be evaluated, rather than as it exists during the designing period of the study.

Also consider the costs involved, e.g. paper and printing, post and phone, travel, and computer software, as well as the time of the personnel.

What methods are best?

The way information is presented can be crucial to how seriously key parties perceive the study. Different types of information convince different people. Equally, the form in which information is gathered restricts the ways in which it can be used. Quantitative measurements and hard facts may be of more use in demonstrating concrete achievement to funders and top management, but qualitative feedback is generally far more useful in establishing improvements necessary for users of a system, or to benefit students on a course.

Resource levels will restrict the amount of information you can most usefully gather and process, but the most sensible method will be dictated by the driving force for the study, accountability, and whether it is intended to be a formative or summative study. The information you choose to gather will ultimately affect the tools and techniques you adopt, with consequences for the resources you require to complete the study successfully.

A key part of the planning is to choose appropriate sources of information (e.g. students, staff, documents) and methods of collecting evidence. Much of this book is designed to help you select suitable approaches. The purposes of the evaluation and the practical features will have some impact on your methodology. Use a variety of methods so that findings from one source can substantiate others. Or the findings from one method can help the design of another, e.g. topics from a group discussion can lead to some of the questions in a survey; comments from the survey could identify issues to be explored in interviews.

It is important to collect as much information as appropriate, but not to exceed the resource base available. The information gathered will need to be refined from one study to the next. Some material will be shown to be less useful than anticipated, while other areas will throw up gaps that would benefit from further examination. Methods of evaluation can also be changed or adapted to fit in with the practicalities of the situation. As each study develops, the process of defining the next study will become progressively easier.

What impact will it have?

Evaluation can be a delicate topic and should be handled sensitively. If you ask similar questions about an innovation of students, lecturers and technicians, you may get conflicting views, so you will need to decide how to cope with the situation. Do not ask questions that raise unrealistic hopes.

How will you support a lecturer who gets a lot of negative comments from students? Some aspects may need to be confidential and anonymous. How will you monitor and deal with unintended outcomes? Many potentially difficult situations can be avoided if you explain the purpose of the evaluation in advance and if you share the outcomes with all involved.

The study is your opportunity to make contact with those people who can provide the best feedback on the area of work being evaluated. Who are the people who will be most affected by your work? Who will use what you are creating? What are their needs? How do you think they might be able to help you? Can you use the study to make contact with external groups by providing a common purpose? What information are you missing that has to be gained from other sources? Naturally, you do not want to alienate any of these groups, so thought about how you approach them will make your evaluation run more smoothly.

What are your deliverables?

How will the results of the study be distributed and to whom? How will the results be implemented into your work? Will responses be directly fed back into course, product, or system, or will a formal report of some type be required? Should you publish the results? Do you perhaps need several forms of presentation depending on the group of people requiring the results?

As you consider each of the above questions, a structure or framework for the study should evolve. This may show that a series of studies would be more valuable. These may be divided into evaluation phases each building on the information generated by the previous phase, or you may design a series of smaller studies, each dealing with a different aspect of knowledge. You must keep the study design flexible to allow for adaptations as results are obtained or as requirements change. The process of evaluation is iterative, and each study must be based on both current needs and previous findings. Working within tight time and resource constraints makes it more and more important to get the initial question right each time.

Gaye Manwaring
Northern College

Gayle Calverley
Academic Services Learning Development
The University of Hull

As with any process, one of the most important steps in carrying out a successful evaluation is choosing the right way to go about doing it. If the study's design is well suited to the questions being considered, the whole process will be made considerably easier.

The surest way to get the design right is through experience. If you have already carried out evaluations, so much the better – you will already be aware of many of the problems and issues that will affect the work, and will be familiar with some of the approaches that can be used to counteract these. However, whether or not you have this experience, there will be other people who do. It is always worth talking to other evaluators when preparing a study, as their perspective will help to identify any potential difficulties in your plans. Further familiarity can be gained through reading about other peoples' studies and approaches. The recipe pages in this book provide an invaluable starting point for this process.

However, it is possible to make sensible choices without needing to become an expert in the topic first. There are a number of questions that can help to choose which methodology is best suited to the topic of your study. These can be grouped into questions about the methodology itself, about the techniques it uses for gathering data, and about how these data are analysed. It should be noted that this approach necessarily relies on generalisations. Whilst they provide a good starting point for evaluation design, practice and experience will obviously enable you to make more informed decisions.

Choosing a methodology

Each methodology represents a different approach to evaluation. The fact that there are so many approaches in common use simply reflects the fact that no single methodology is 'the best'. Which one will be most appropriate for you depends on the type of questions you are asking. It's important to be clear what your questions are – apparently subtle changes can have considerable impact. Asking, "What factors influence how well students do?" suggests an exploratory study, which seeks to identify influences on performance. However, asking "Which of the following factors influences how well students do?" suggests a comparative study, possibly involving a controlled experiment.

With a clear question in mind, it is possible to start working out which methodology you need. A good starting point is to decide how exploratory your study needs to be. In the example above, the 'what' question is highly exploratory – the evaluator

has little or no idea about the factors that will influence learning. These need to be discovered in the course of the study. In the 'which' question, factors have already been found. What remains is to test them to demonstrate their influence. For this reason, the study needs to be much less explorative. Open, qualitative methodologies such as interviews, observations and concept maps tend to be best suited to explorative studies, whilst checklists and experiments require a framework for questions to be fixed in advance.

A second important question to ask is how authentic your study needs to be. When designing instructional material, it may be more appropriate (and more ethical) to test your ideas in a laboratory-like setting, rather than on students whose exam grades may be affected. However, such controlled setups are unsuitable for evaluating how to improve the way that these materials are used as part of the curriculum. Such studies require a more authentic setting. Clearly, controlled experiments are far less authentic than (for example) ethnographic studies or student profiling. Some techniques, however, can be used in both types of setting – observations are a good example of this.

Finally, it is important to be aware that the number of people who will be involved in the study will have an impact on the approach you choose. It would be impractical to carry out open interviews with 200 students, and probably inappropriate to use a comparative experimental design on a group of eight participants. Broadly speaking, the methodologies that are best suited to large groups will limit the amount of qualitative data to be gathered.

Gathering data

Just as it is important to be clear what questions are being asked, it is also vital to be clear about what will count as evidence. When asking about students' performance, are attitudes and perceptions important? What about the way in which they interact with the instructional material, or with each other? Do you view performance on tests or in exams to be the same as learning or as understanding of key concepts? If not, what can demonstrate this?

One characteristic of evaluation methodologies is the types and range of data that are collected. As with methodologies, the process of choosing the right data capture techniques can be made easier by considering a series of questions. Perhaps the simplest to ask is how objective the data that is gathered will be. If subjective information, such as attitudes and perceptions, are of interest then questionnaires and interviews are appropriate. If you need to know how students act or interact, or how their performance is affected by some change in the curriculum, data such as those contained in video logs or test results will be important. It is worth emphasising that the subjectivity of data has no impact on the rigour with which it will be analysed – it is simply a description of the type of data under consideration.

Another important quality is how focused the data will be. One reason that ethnographic studies use observations is that peripheral data, such as weather, social interactions and so on can all be taken into account if they seem to influence proceedings. A multiple-choice questionnaire, on the other hand, gives no latitude in the information that is gathered. The tight focus will make the data easier to analyse, but the down-side to this simplicity is that it means that data are also limited.

There is the practical concern of how long data gathering will take. Participant observation is extremely time intensive, as are interviews. Video recordings and questionnaires simply require setting up and gathering in, making far fewer demands on the evaluator's time.

Finally, it is worth being aware that access to resources can restrict the range of data capture techniques that can be used. The availability of screen capture software, video cameras, microphones and even an adequate photocopying budget will all need to be taken into account.

Data analysis

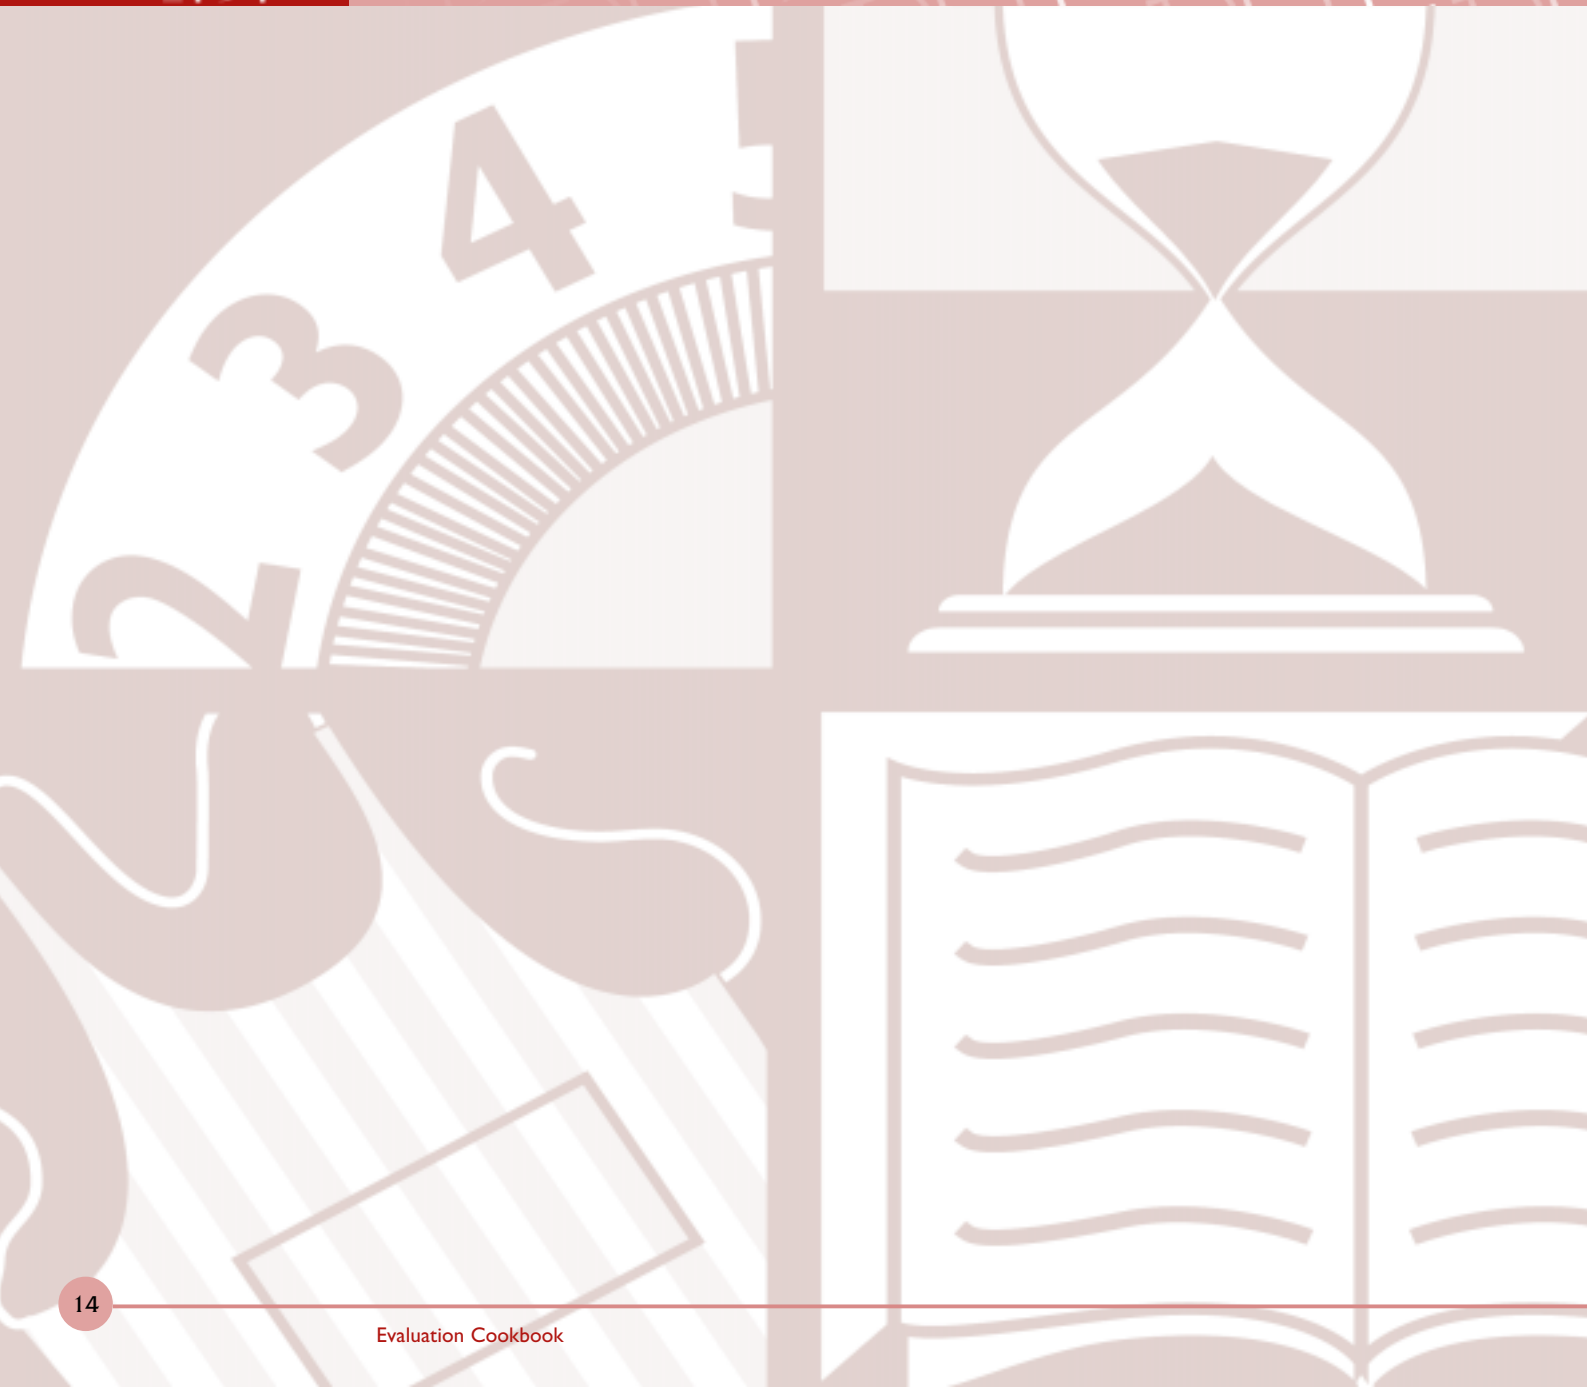
As with gathering the data, an important concern here is how long the data will take to analyse. The process of analysing and categorising qualitative data can be very time intensive. Transcription (which could be considered to be part of either data capture or analysis) also needs to be considered. As a rule of thumb, transcribing one hour's audio tape can take from two to four hours; one hour of video can take from six to eight, depending on how practised the transcriber is and how thorough the final account needs to be. By contrast, with the help of software packages, descriptive and inferential statistics can be dealt with very quickly. In such situations, it will often take longer to design a controlled experiment than it will to analyse the results.

Finally, it is worth being aware that data can be presented in a range of formats, which will be appropriate for a different purpose or audience. These formats will be restricted by the type of data gathered and the methods of analysis. Qualitative data is extremely good for presenting illustrative or personalised information. It is less useful for providing summaries or overviews, however, and unlike analytical statistics, it is hard to specify how confident you are that the findings have been caused by a particular factor, or that they will generalise.

Summary

Although it's important to choose the right methodology, and there is a wide range of approaches to choose from, making the choice does not need to be daunting. Once the evaluation question has been chosen, it is possible to work out the characteristics that a suitable methodology needs to have. Armed with this description, all that remains is to review the options that are open to you and choose the one that meets your needs.

Martin Oliver & Grainne Conole
LaTID, The Learning Centre,
University of North London





prep time



time/student



time/admin








analysis



additional resources

This aims to provide a rough guide to the relative resourcing implications involved in carrying out each different evaluation method. However, it is important to note that these times are only approximations and the actual times will depend on a range of factors such as the number of students, the level of resourcing within your institution and your experience in carrying out evaluations.

					
Checklists	low-moderate	low	low	low	low
Concept maps	low	low	low	low	low
Confidence logs	low-moderate	low	low	moderate	none
Cost effectiveness	moderate-high	none	none	moderate-high	none
Designing experiments	high	low-moderate	low-moderate	low	low
Ethnography	low	low	high	high	moderate
Focus groups	low	moderate	moderate	low-moderate	moderate
Interviews	moderate-high	moderate	high	moderate-high	moderate
Nominal group techniques	low	low	low	low	low
Pre and post testing	high	moderate	moderate-high	moderate	low
Questionnaires	moderate	low	low	moderate	none
Resource questionnaires	low	low	low	moderate	none
Split screen video	moderate	low	moderate	moderate	high
Supplemental observation	low-moderate	moderate	moderate	moderate	moderate
System log data	moderate-high	low	low	moderate	moderate
Trials	moderate	high	moderate	moderate	low-high

key:



preparation time – The time taken for planning and organising the evaluation



time/student – The amount of your student's time needed for the evaluation



time/administration – The time needed to conduct the evaluation



analysis – The time taken to analyse the data



additional resource – Any additional resources needed as a direct result of the evaluation

Colin Milligan
TALISMAN
Heriot-Watt University.

A quick way of getting a lot of information about an implementation: data gathering is efficient and collection and analysis can be automated. Although the data gathered is rather low level there can be lots of it.

Uses

- ◆ Checking whether a range of materials fit certain standard criteria
- ◆ Measuring how well you have matched your materials to your students' needs
- ◆ Obtaining feedback from large population samples

Process

1. Planning

Make sure a checklist is appropriate for the type of analysis you wish to carry out; on the one hand, checklists are good for quickly identifying issues, but they may not provide enough information to allow you to rectify any problems.

Look at the material that you want to analyse with the checklist; did you have objectives which you can specifically test to see if they have been met? Can you do the same with previously unstated objectives?

2. Identifying your population sample

Decide who your population is and whether there is any information about their previous experience – for example, qualifications, previous courses, expectations etc. – which may help you interpret the information they provide. You can ask for this information in the checklist.

3. Design

Carefully choose the best question type. Often, you may want a simple yes/no answer, e.g. did you find **X** useful, was it easy to carry out **Y**, etc. However, sometimes supplementary choices are appropriate. For instance, when assessing whether objectives have been met it might be useful to determine whether the respondents felt the objectives had been fully or partly met. Whether respondents had prior knowledge of the material might also modify the meaning of their answer.

As closed questions are easy to answer you can ask many questions at once without risking overloading the user.

Keep the wording clear, trying not to introduce terminology. Rather, try to directly relate the question to specific parts of the materials, such as objectives.

Try to group questions logically. Use subheadings and clear instructions to lead the users through the questions.

Pilot the checklist with someone who knows the material. As well as comments on clarity etc., they may be able to comment on the balance of the questions.

Example Checklist

Which of the following elements of the course did you find useful? Please tick all that apply:

- | Week One | Week Two | General |
|--|--|---------------------------------------|
| <input type="checkbox"/> Course text | <input type="checkbox"/> Course text | <input type="checkbox"/> Submit a URL |
| <input type="checkbox"/> Dialogue discussion | <input type="checkbox"/> Links | |
| <input type="checkbox"/> Digest | <input type="checkbox"/> Mailing list | |
| | <input type="checkbox"/> Dialogue discussion | |

Use checklists to give you quick feedback on a new implementation.

When using the same materials with different groups of students, an evaluation checklist can quickly highlight whether the needs of one group are not being properly met.

Asking the same question in more than one way can help to reduce ambiguity in the final analysis.

Vary the wording of questions so that respondents aren't always providing the same response. They'll get bored and won't pay due attention.

If you set aside class time for completion of the checklist, it is more likely to get done; otherwise your students will find something more pressing to do.

4. Delivery

Paper or electronic delivery of the checklist evaluation is possible. WWW based forms can be used to collect responses efficiently, but should only be used where appropriate – for instance when the materials being evaluated have already been delivered electronically or possibly when face-to-face collecting of checklist feedback is not practical.

Low response-rate is a considerable problem with checklists as with other form-filling evaluations. You might want to consider how you can make completion of the checklist more relevant to the students - by making them feel that they will get something out of the exercise. For instance, you may make the questions relevant to revision by reiterating the objectives and asking the students to indicate whether or not they felt they had been met. You could also include ideas for further study/reading for each objective.

Time your evaluation carefully – should your checklist be delivered directly after the implementation, whilst the class mind is focused (and the whole class is captive) or after a period of time (when concepts have been strengthened but individual details lost)?

For electronic forms, values can be assigned to responses, enabling automation of the tallying process. Electronic submissions can often be formatted for direct import into an appropriate analysis program.

5. Analysis

In addition to a collective analysis i.e. what proportion felt that a particular objective had not been met, you may want to relate different answers from the same respondent. Alternatively you could group students according to their responses to one particular question.

A checklist can often be used very effectively as one component in an evaluation – possibly to identify specific issues that can be investigated further in a focus group or structured interview.

Try to feedback your results to your students and to follow-up any recommendations.

Variation

Checklists can also be used by lecturers while selecting resources to enhance teaching. Use a list of attributes that you think are required for a successful implementation to provide some guidance when looking at new software. Attributes might include: the software is cheap, the subject content is accurate, or the software engages the user in activities which are relevant to the learning objectives.

Other Relevant Pages

Recipes

- Designing experiments
- Split screen video
- Cost effectiveness
- Trials
- Pre and post testing
- Resource questionnaires
- Interviews
- Focus groups

Information Pages

- Likert scales
- Statistics questions
- Pre and post testing
- Questionnaires
- Guidelines for questionnaires
- Student sample

Serving suggestions

References

How to Evaluate Learning with Technology?

Recent developments in CAL evaluation methodology show a definite shift from isolation to integration as the interdependence of content, context and individual characteristics became increasingly recognized. This shift reflects concurrent developments in learning research and evaluation in non-technology related fields. Although attempts to separate the strands seem redundant as technology continues to permeate all functions of higher education institutions. While the reasons for evaluation may remain the same, i.e. :

- ◆ to assess (and improve) the effectiveness of whole courses and their various components
- ◆ to identify the influences and effects of various contextual factors

the rationale, assumptions and methods have changed considerably during the relatively short history of the discipline.

When computer assisted learning first became popular in the 1960s, evaluation typically meant attempting to isolate the effects of a single resource, application of sampling methods designed to balance individual differences among the study population and creation of a 'clean' experimental situation where objective truth about the impact of a particular intervention could be revealed. Thankfully, CAL technology is not the only thing that has come a long way since the 60s. Learning evaluation as a discipline, and studies of CAL in particular, have developed, through experience, into something infinitely more sensitive to the impact of innovations and appreciative of the influence of personal and contextual factors such as prior knowledge, learning style, integration into course structures, instructional strategy, design and support. In fact, the basis has shifted through 180 degrees, from a predictive, hypothesis testing model to a responsive process from which hypothesis or theory generation is the outcome. A brief and approximate history of developments reveals the following milestone events.

1960s

CAL Types

Computer assisted instruction, programmed learning, branching programs

Evaluation

Controlled, experimental studies based on the behaviourist, measurement oriented paradigm articulated in the 1930s by Ralph Tyler and Skinnerian stimulus – response related assumptions about learning. Learning is still regarded as independent of subject or context.

Findings

Scores and outcomes based, no relevance attached to process or contextual factors

References

- Robson, C. (1993). *Real World Research*. Oxford UK and Cambridge USA: Blackwell.
- Skinner, B. F. (1954). *The Science of Learning and the Art of Teaching*. *Harvard Educational Review*, XXIV(2), 86-97.
- Skinner, B. F. (1968). *The Technology of Teaching*: Appleton-Century-Crofts.
- Tyler, R. (1950). *Basic Principles of Curriculum and Instruction*.

1970s

CAL Types

Tutorial programs, simulations.

Evaluation

Still predominantly experimental but with an emerging counter-culture, traceable to the 60s, that argued for process oriented descriptions of programs in use in specific situations, and recognized the importance of social, political and economic factors. Methods associated with the 'new' evaluation are varied and include interviews, questionnaires, profiles, think aloud protocols, observations etc.

Findings

Descriptive and indicative of many contributory factors to effective learning outcomes, e.g. teaching and learning styles, prior knowledge, motivation, classroom culture, assessment. Initially case specific though generalizable through grounded theory type development.

References

- Laurillard, D. (1978). Evaluation of Student Learning in CAL. *Computers and Education*, 2, 259-265.
- MacDonald, B., & Jenkins, D. (1979). *Understanding Computer Assisted Learning*. Norwich: University of East Anglia.
- Parlett, M., & Hamilton D. (1976). *Evaluation as Illumination: A New Approach to the Study of Innovatory Programs*. In D. A. Tawney (Ed.), *Curriculum Evaluation Today: Trends and Implications*. London: Macmillan Education.
- Stake, R. (1975). *Evaluating the Arts in Education, a Responsive Approach*. Ohio: Columbus.

1980s

CAL Types

Microworlds, complex simulations, intelligent tutoring, generative programs

Evaluation

The need for responsive/evaluative methods is clear but academic credibility for the qualitative methodology is still hard won. Naturalistic methods based on the interpretive and critical paradigms are increasingly popular as experimental methods consistently fail to produce sufficient detail for designers' and evaluators purposes in formative and summative studies. Usability studies take precedence over learning evaluation and CAL design guidelines and standards evolve.

Findings

Results of formative evaluation and various forms of user testing become important inputs to development, and the iterative design cycle is established. Case and situation specific factors are identified and reported as the shift away from large experimental studies and generalizable results on learning issues continues.

References

- Bates, T. (1981). Towards a Better Research Framework for Evaluating the Effectiveness of Educational Media. *British Journal of Educational Technology*, 12(3), 215-233.
- Marton, F. (1981). *Phenomenography - Describing Conceptions of the World Around Us*. *Instructional Science*, 10, 177-200.
- Marton, F. (1987). *Phenomenography - A Research Approach to Investigating Different Understandings of Reality*.
- Reeves, T. C., & Lent, R. M. (1984). *Levels of Evaluation for Computer-Based Instruction*. In D. E. Walker & R. D. Hess (Eds.), *Instructional Software : Principles and Perspectives for Design and Use*. Belmont, CA: Wadsworth Publishing.
- Walker, D. F., & Hess, R. D. (1984). *Evaluation in Courseware Development*. In D. F. Walker & R. D. Hess (Eds.), *Instructional Software : Principles and Perspectives for Design and Use*. Belmont, CA: Wadsworth Publishing.

1990s

CAL Types

Online courses, user generated resources, full multimedia simulations and tutorial CAL

Evaluation

Integrative response studies are conducted in authentic contexts using mixed methods and multiple data sources. Methods must accommodate situations where teachers and learners may never meet face to face. Evaluation is now accepted as an important and ongoing aspect of program and course improvement, the importance of context is undisputed and attempts to isolate the effects of CAL are less relevant than assessment of how it works in conjunction with other resources.

Findings

Part of an ongoing process which feeds back into a plan - implement - evaluate - improve loop. Learning objectives, means of assessment and opportunities for data collection are determinants of what findings will be sought and how they will be used. Studies involve qualitative and quantitative measures as appropriate.

References

- 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 Glasgow: TILT Project, 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 & Education*, 26(1-3), 17-32.
- Ehrmann, S. (1995). Asking the Right Questions : What Does Research Tell Us About Technology and Higher Learning? *Change Magazine* (March/April).
- Gunn, C. (1996). CAL Evaluation : What Questions Are Being Answered? *Computers & Education*, 27(4), 57-60.
- Keane, D. R., Norman, G. R., & Vickers, J. (1991). The Inadequacy of Recent Research on Computer Assisted Instruction. *Academic Medicine*, 66(8), 44-48.
- Kulik, C. (1991). Effectiveness of Computer Based Instruction: An Updated Analysis. *Computers in Human Behaviour*, vol 7 (1-2), 75-95.
- Milne, J., & Heath, S. (1997). *Evaluation Handbook for Successful Courseware Development*, Centre for CBL in Land Use and Environmental Sciences, Aberdeen University.
- Reeves, T. C. (1991). Ten Commandments for the Evaluation of Interactive Multimedia in Higher Education. *Journal of Computing in Higher Education*, 2(2), 84-113.
- Somekh, B. (1990). The Evaluation of Teaching with Computers. *CTISS File*, 10, 32-39.
- Wills, S., & McNaught, C. (1996). Evaluation of Computer Based Learning in Higher Education. *Journal of Computing in Higher Education*, 7(2), 106-128.

Conclusions

A number of valid approaches to evaluation are currently in use: one common variation being in how broadly the term is defined. A narrow perspective is where the effectiveness of a particular program or part of a program is assessed in relative isolation from the wider context in which it is used. An example of this would be where a tutorial program for teaching the economics concept of price is evaluated immediately following students use of the program. Demonstrated understanding of the concept would be one measure of effectiveness, ability to apply it in different situations may be another. It would be useful to know, e.g. if students had any prior knowledge of the concept, had learned it from a textbook or other source, then reinforced it through use of the CAL program, and whether they would be able to transfer the new concept to other applicable subjects such as accounting or marketing. A broader perspective might include how well the CAL program is integrated into the whole course and assessment structure, and how CAL use in general is viewed by students, presented by lecturers and supported by the institution. All these factors can influence the effectiveness of learning outcomes, even although they may not relate directly to the design and use of a particular piece of courseware.

It may be concluded then, that the purpose of the evaluation will define its scope. Courseware developers may be more concerned with the design related aspects while organization, policy or staff developers may tend to look at the broader picture. However, all perspectives require some attention to contextual factors and the influence they bring to students use of courseware and the effectiveness, or otherwise, of learning outcomes.

Cathy Gunn

Education Technologies Advisor,
CPD, University of Auckland.

A typical question using a Likert Scale might pose a statement and ask the respondent whether they Strongly Agree - Agree - Undecided - Disagree or Strongly Disagree.

The responses elicited may be coded e.g. 1-2-3-4-5, but this remains just a coding. It makes no sense to add a response of agree (coded as 2) to a response of undecided (coded as 3) to get a 'mean' response of 2.5 (what would it mean?). So how **can** you analyse data from a Likert scale?

The data collected are ordinal: they have an inherent order or sequence, but one cannot assume that the respondent means that the difference between agreeing and strongly agreeing is the same as between agreeing and being undecided.

Descriptive Techniques

- ◆ Summarise using a median or a mode (not a mean); the mode is probably the most suitable for easy interpretation.
- ◆ Express variability in terms of the range or inter quartile range (not the standard deviation).
- ◆ Display the distribution of observations in a dotplot or a barchart (it can't be a histogram, because the data is not continuous).

Inferential Techniques


Investigate differences between (the medians of) comparable groups using non-parametric methods e.g.:

- ◆ for two unrelated samples – Mann Whitney test;
- ◆ for paired samples – Wilcoxon signed rank test;
- ◆ for three or more samples – Kruskal Wallis test.

Investigate any association between two different sets of observations using a chi-squared test of association.


Nora Moge
Coordinator, LTDI,
Heriot-Watt University.

Bar chart showing the number of students in different faculties



- Features**
- Bars are separate
 - can be used for nominal or ordinal data
 - only the height of the bar is relevant, not the width

Histogram of students reaction times



- Features**
- bars must touch
 - can have bars of different width – but need to adjust their height accordingly
 - the area of each bar represents the number of observations

Judith George
The Open University, Scotland.

A concept map or a mind map is a visual representation of the links or associations between different concepts or pieces of information.

Uses

- ◆ At the end of a tutorial, in the last 10 minutes, to consolidate learning, check understanding
- ◆ At the start of a session, to map where students are

Process involved

1. Use an example

For the first time, have one you prepared earlier, in case the students are not familiar with concept maps. Students can be quite unused to doing this - so do be encouraging and ready to be the guinea pig yourself.

2. Draw the concept maps

Ask everyone to draw their own map of what was significant for them in the session – the main ideas, facts and so on. Draw one yourself.

3. Compare Maps

The first time show your version, and ask if it bears any resemblance to theirs. Some bold soul will venture theirs - and then the rest.

4. Reflect on the Maps

Look for feedback about coverage - and gaps, or misunderstandings. Encourage them to articulate and explain their schemes.

Briefly confirm important points which have been well understood, and correct any minor misunderstandings. If there are larger areas of doubt or misunderstanding, put those on the agenda for next time.

Variations on this technique

Concept/spider/mind maps can be used for many purposes; for planning written work or presentations, for self-review of understanding, for feedback. The more familiar and fluent students are with this technique, the more they will benefit from any application of it.

Try selling the method, which takes valuable class time, as an experiment which you will abandon, if they are not convinced that it pays off in consolidating and checking learning.

Don't expect great things the first time.

Be positive and encouraging; build on positives; use the opportunity to get further discussion of how you handled points which they found difficult: - "Was it the examples I used? Could we have tackled this more helpfully?"



low



low



low




low



low

Example of a Concept Map



Other Relevant Pages

Recipes

- Focus groups
- Nominal group technique
- Confidence logs

Information Pages

- Learning styles
- Pedagogic toolkit
- Working with groups
- Isolation or integration
- Pre and post testing

Serving suggestions

References

Students can be worried about the time this takes within a class. Sell it with confidence on the basis that it will actually get them further forward than a final alternative 10 minutes of input would - their actual experience will confirm this. It may well also open up dialogue between you and them about the process and agenda of your sessions, deepening their motivation and commitment.

CAL packages have come a long way from a set of floppies with merely text and graphics on them, to highly sophisticated 'interactive' learning modules. Often these packages claim learning is made easy for students, by taking their learning styles into account.

But is this truly the case? Find below some learning styles explained and translated into good practice for use in CAL packages.

Auditive or visual learning styles

One of the easiest recognisable learning styles that students have is the preference for auditive intake or visual intake of taught knowledge. Students usually either prefer look at material individually, or prefer to listen to lectures. Ideally a CAL package caters for both. A module supported with audio as well as a visual representation of content is easily achievable using modern multimedia facilities.

Supporting study styles

Tutors often explain to students how they should write essays, how to prepare for exams, how to deal with big amounts of reading work, i.e. how to study. In doing this, tutors encourage students to use better methods of learning. The better the lecture, the better this advice is ingrained in their teaching.

However, very few CAL packages support this students' need. Occasionally, a button featuring 'study advice' is available, and although this is a fairly artificial solution, it is a good start. Weaving study-related remarks into the body of a module within a package would be ideal.

Allowing for failure

In a similar way to toddlers learning how to walk by trial and error, students need opportunities to stumble and fall within their learning process. This helps them to develop knowledge that they can use creatively. However, some packages do not allow for failure. Modules are immediately followed by assessments, without allowing for a 'test-area' or playground. It is good practice to leave learning space for package users.

Interaction: none of it, false or true interaction?

Broadly speaking, all CAL packages try to incorporate interaction. Often this is explained by developers as 'the package giving immediate feedback to students input'. Quite often the feedback is no more than 'Right', 'Wrong' or 'Please try again!'. This is false interaction and serves very little pedagogical purpose. Usually students will resort to a trial and error approach to this sort of interaction, with all attention going to the result of their actions, not the content.

Pedagogically successful packages give an explanation of why something is right or wrong and preferably do not allow for retrying based on trial and error. True interaction is concerned when a student's input steers the activities of the package. This could, for instance, mean that a students' right answer allows him to skip the next module, or a wrong answer adds a remedial module.

Please note that even in a mass-lecture where students are not encouraged to ask questions or discuss topics, learning often takes place in a two-step fashion. The first step is passive: the student listens to what the lecturer talks about, reads slides or looks at the blackboard. The second step is active: students generally make notes and therefore engage with the material they need to learn, however shallow this may be.

Students observed while using CAL packages showed very little active learning, unless the package promotes interaction. They are encouraged to work together or guidance is given on a more active way of working through a package (for instance by making notes).

Separation of knowledge systems.

Human beings can only retain knowledge by linking new knowledge to existing knowledge. To make factual knowledge easily accessible, it is important that one fact is linked with as many other facts as possible. It is a bit like in a cardfile: by introducing as many crosslinks as possible, any item is easy to find.

It is easiest for students to pick up knowledge if it relates to something they already know. By referring to as many topics as possible, which have nothing or very little to do with the content of the CAL package, effective students' learning can be supported. Of course this can also be done by the tutor or lecturer teaching the course, if the package is only part of a full course.

However, for developers of CAL packages it is important not to let students develop a system of knowledge solely related to the CAL package concerned, but link knowledge to a wider area within the course.

(As an illustration for tutors/lecturers: have you ever wondered why students do brilliantly at your exams, but six months later, they cannot use any of the skills you have taught them in other courses? You may be looking at the effect of separate knowledge systems).

Gwen van der Velden
Quality Promotion Officer,
University of Kent

Helyn Thornbury
University of Strathclyde.

These are self assessment measures which are used to gauge a student's confidence level in a particular part of a course.

Uses

- ◆ Usually used in conjunction with other methods
- ◆ Gives an indication of the learning development of students as a comparative measure (before and after an activity)
- ◆ Provides a 'snapshot' of the class at a given point

Process

1. Identify situation to investigate

Confidence logs should be applied to identified situations. These can be:

- ◆ the particular point(s) in the class at which you wish to investigate the students' development – snapshot(s);
 - ◆ an activity to investigate e.g. a scheduled lab session using a computer simulated experiment – comparative study.
- For a comparative study, note that you will need to apply the logs before the identified situation as well as after.

2. Identify key areas of interest

You may be seeking information that is either general or specific.

General – Relating to the knowledge/skills objectives of the activity/course.

Specific – Relating to areas of interest (e.g. where you suspect there may be an area of weakness in the material).

In either case, define these interests before designing the statements for the log.

3. Construct statements

Many of the issues relating to the construction of questions for a questionnaire are equally valid in the context of confidence logs. However, it is important to remember the statements are not questions and should not be constructed as such.

Generally, they should be short, concise and unambiguous. Keep each statement simple; avoid combining areas together in one statement (or using high level statements) – if the student feels very confident in one aspect and not in the other it is very difficult for them to rate their confidence and fairly meaningless for you.

Remember to use language which the students will understand, particularly if you think they may have experience of the material in another context. Also, remember that there is a difference between an understanding of a theoretical and a technical competency so be clear what you are interested in, and communicate this to the student! (See box at end.)

Example Confidence Log

Topic	very confident	confident	some confidence	little confidence	no confidence
Calculating long division		✓			

Think about practical issues such as access and time available.

Don't use more than 12 statements – keep the log short for the student to fill in.

The knowledge and skills objectives of an activity/course may have previously been constructed in a useful form for Teaching Quality Assessment (TQA) specified documentation.

Remember to schedule time for the students to fill in the logs - or you may lose data when they run out of time and leave for their next class!

Tell the students that no knowledge of something equals no confidence! Otherwise they sometimes leave it blank!

4. Practicalities of application

There are various practical issues to consider. They can be summarised as when, where and how.

When

Comparative – For comparative studies, the when will be before and after a specific activity. Some issues to consider are: Is the activity timetabled? If it isn't, do you know when the students will do the activity? Do all the students work together or are they split into groups?

Snapshots – For a snapshot study, the when is related not to a specific activity but to a point in the class. As well as issues similar to those addressed for a comparative study, you should also consider whether or not all the students are expected to reach the specified point at the same time.

Where will you be able to get access to the students? In an existing lecture or lab., or will you need to arrange something?

How will the logs be presented? Part of the how is the introduction/instructions you give to the students. Some people use a short introductory paragraph to explain what is required of the student, others give these instructions verbally, some use both. This is really dependant on the situation but the students should receive some guidance. Other aspects you might consider include: whether the logs are going to be electronic- or paper-based; who, if anyone, will need to be there to administer the logs; and how long will it take to complete the logs?

5. Analysis

There are many ways to visually present the information from logs. Bar charts can be recommended as a simple and informative option.

Comparative – If you have anonymous logs you can look at the spread of confidence in a standard bar chart format for each of the statements. Comparing the bar charts before and after an activity will give you a general indication of any shift. If you can pair before and after confidence logs, this will allow you to chart changes in confidence which is generally more informative.

Snapshots – By constructing bar charts for each statement (as above) you can gain an overall impression of the confidence of the class at a given moment which can be compared with your expectations.

Other Relevant Pages:

Recipes

- Questionnaires
- Checklists
- Pre and post testing
- Focus groups
- Designing experiments

Information pages

- Statistics questions
- Likert scales
- Learning styles
- Student sample
- Guidelines for questions
- Pre and post tests
- Pedagogic toolkit

Serving Suggestions

References

Variation

Confidence logs can also be used longitudinally – in the same way as the snapshot outlined above but repeated over a period of time. This can allow you to look at the development of confidence over your whole class. By examining the logs you can check whether the development in the different areas matches your expectations, and you can look for groupings in the students. However, this type of analysis involves an associated increase in time and, as a consequence of the repetition, can negatively affect the quality of data.

Theory v Practice. Particularly in very practical areas, there can be a difference between an understanding from a theoretical perspective and a practical competency. If an activity is a combination of theoretical understanding and development of skill, separate out the two aspects in the logs.

Philip Crompton
University of Stirling.

Any educational intervention has an outcome and a cost. A measure of the cost effectiveness is obtained by measuring COSTS against OUTCOMES.

Uses

Analysing cost effectiveness of an intervention can involve evaluating various options, for example:

- ◆ reviewing a number of alternatives all of which are within the realm of cost feasibility;
- ◆ considering which supplemental programs ought to be used to improve an educational outcome;
- ◆ trying to identify which program has the best average outcome per student relative to the per-student cost.

Process

COSTS

1. Identify all the resources

Identify all the resources necessary to create or replicate the intervention and its observable effect(s) even those not included in budgeting expenditures e.g. materials, operating costs or general maintenance of equipment and resources.

2. Categorise the resources

Categorise each of the resources under various headings e.g. personnel, facilities, equipment, client inputs etc. This facilitates comparison between various alternatives within an intervention.

3. Calculate the overall costs of the intervention

OUTCOMES

4. Identify the outcomes for analysis.

Outcomes are the results of the intervention, for example, higher student motivation, improved student performance in assessments, lower student drop out rates.

5. Carry out the analysis

Analyse on the basis of your selected COSTS and OUTCOMES. Is the outcome worth the costs?



Try and specify each of these resources in order that their value can be ascertained.

Headings under which the resources are placed should be kept consistent throughout.

The measurement of effectiveness is determined by the objective selected for analysis.

EXAMPLES OF EFFECTIVENESS MEASURES

Program Objective	Measure of Effectiveness
Reducing dropouts	Number of potential dropouts who graduate.
Program completions	Number of students completing a program.
Student learning	Test scores in appropriate domains utilising appropriate test instruments.
Employment of graduates	Number of graduates placed in appropriate jobs.
Student satisfaction	Student assessment of program on appropriate instrument to measure satisfaction.
Physical performance	Evaluations of student physical condition and physical skills.

A number of arguments have placed both for and against carrying out cost effectiveness exercises:

Against:

- ◆ since educational technology is a tool that staff and students use, any evaluation of its educational effectiveness (or lack thereof) is necessarily a critique of those staff and students,
- ◆ why take on the risk and cost of evaluating educational impact – can you afford the answers?
- ◆ just search through educational research, perhaps other studies will reveal if it is cost-effective,

For:

- ◆ even if we find evidence in professional literature it doesn't tell us about how and if **our** investments have been worth the money;
- ◆ budgets are politically vulnerable, if we don't provide educational evidence to reshape or defend our budget, it can be (irrationally) cut.

"Cost-effectiveness analyses, can be (and are) manipulated ... basically these cost ratios can be as low as we want."

(Fletcher, 1990)

Other Relevant Pages

Recipes

- Resource questionnaires
- Interviews
- Trials
- Focus groups
- Pre and post testing
- Designing experiments
- Nominal group technique
- Ethnography

Information Pages

- Pedagogical toolkit
- Student learning
- Pre and post testing
- Isolation or integration

References

Questions to ask yourself before you invest too much effort in analysing your data.

1. What will a statistical analysis tell you that you couldn't establish in other ways?
2. Have you thought about your sample and how it was selected? Are the number of missing responses satisfactory? Why are they missing?
3. Does your data include outliers or unusual observations? Why? Will they have an undue influence on your conclusions?
4. Is the size of your sample sufficient to make a full analysis worthwhile?
5. Have you used graphical methods to explore your data and get a good feel for it?
6. Are there any confounding variables?
7. Do you know what type of data you have collected? (Nominal? Ordinal? Interval? Ratio?)
8. Have you considered using confidence intervals rather than/in addition to an hypothesis test?
9. Do you know what conditions are assumed for the test you are considering to be valid? Do your data meet those conditions?
10. Just how reliable are your data anyway?

Nora Moge
Coordinator, LTDI,
Heriot-Watt University.

Carrying out an evaluation involves time for both you and your students. Whether it is time taken out of a lecture or during a lunch hour, it is time away from a student's study or social time. Therefore, before you start, it is important to be clear in what you want to find out from your evaluation, which evaluation method is going to be the most appropriate for your needs and how many students it is going to be appropriate and practical to include in your study.

How you go about selecting your student sample will have an effect both on the information gathered and the impact that your findings might have. A number of sampling options are possible but it is worth bearing in mind what the benefits and limitations of each might be.

Letting your sample select itself:

Making materials available for students to look at or handing out questionnaires for students to complete in their own time can result in a disappointing if not unrepresentative sample for an evaluation study.

However, you can gain a snapshot impression of students' general opinions. Watch that you don't try to draw any major conclusions from the responses of a small section of a class sampled in this way. Individuals who just happened to turn up at the last lecture of the term can have very different reactions to the majority of the class with whom you are planning to use a computer package during the following year.

Asking for volunteers:

Again, this can produce an unrepresentative sample for your evaluation. Volunteers are likely to be the most conscientious of the class or the students who are just trying to please you.

However, when you are carrying out a fairly extensive and time consuming evaluation study, you are probably going to depend on the good will of people who are motivated enough to volunteer to help.

Selecting the sample yourself:

If you pick your own sample of students, you have the opportunity of being able to identify the students who are likely to be most co-operative or a group of students with the most appropriate skill levels. You can also select a random sample of students in order to try and get a more representative cross section from the class.

Watch, however, that by selecting one group from a class and involving them in an evaluation study you are not perceived as giving one group of students additional / better / preferential support or tutoring than the rest of the class. It can be easy for students complain that they feel disadvantaged from their peer group in some way.

Involving the whole class in the evaluation study:

This will provide a more representative sample than by taking just a small section, but you could have problems with managing and working with data from an evaluation study with large numbers. If you are planning to involve the whole class, you might also want to consider whether or not you are going to timetable this study as part of the normal syllabus or to include it as an additional extra. Ten minutes to complete a questionnaire at the beginning of the class doesn't cause too much of an upheaval, but asking all of a large class of students to work their way through a CAL package without evaluating the materials with a small group of students beforehand, could prove to be rather an unpredictable exercise.

Jen Harvey
Implementation Support Consultant,
LTDI, Heriot-Watt University.

Hamish Macleod
University of Edinburgh.

Uses

Every educational innovation is an experiment in some sense of the word; you change something about the students' experience, predicting that better learning will take place.

A controlled experiment is a way of teasing out the details of just which aspects of your innovation are influencing the outcomes you are considering and bringing about the changes you observe. The experimental method is a way of thinking about the evaluation process such that all the possible sources of influence are kept in mind.

Process

1. Describe the Innovation

Construct a clear description of the innovation that you are seeking to monitor.

Exactly what will be different in the students' experience after the change you propose as compared to the current situation? The ideal experiment manipulates only one factor at a time, thus enabling very direct causal links to be explored. In practice, a number of changes may have to take place at once for reasons of expedience.

2. Decide the parameters of your experimental design

What are you going to compare with what? Will it involve a comparison of what happened last year (before the initiative) with the experience of the current year (in which the initiative is in place)? Perhaps only part of the class will experience the new learning situation, and their performance (or their judgement of their enjoyment of the course, or whatever) will be compared with that of their immediate colleagues who have not experienced the change. Or perhaps you plan to continue with your normal practice and compare the learning outcomes of your students with those of an equivalent class taught by a colleague at another institution where some sort of innovation has been put in place.

3. Define "success"

Decide what outcome would be needed for you to consider your experiment to be a success. Perhaps the objective is to address some issue of recruitment onto a subsequent level course. Are changes going to be reflected in the students' academic performance? Will they be expected to enjoy their learning experience more, or to express more confidence and satisfaction with their achievements?

4. Decide how to measure successfulness

Decide how your predicted outcome can best be measured. Other sections in this guide have more to say about the different ways in which you can obtain qualitative or quantitative estimates of some dimension which tells you something about the outcomes of interest.

Be aware that what you measure, and what you are interested in, may be subtly or even profoundly different. Some things may be easily measured (like the scores in a multiple-choice examination) while others (like the depth of understanding of some concept) may be considerably more difficult to measure; and the temptation is always to take the simple course. On the other hand, good, simple proxy measures can often be found for the outcome of interest. It is not necessary that the measurement you collect be direct, but only that it is strongly correlated with what you need to know about.

Don't tinker with too many aspects of your course at one time, or it will be impossible to identify just which of these changes you have made has caused the improvement (or disruption) which you observe.

As far as possible, compare like with like, trying to exclude the possibility that any differences you observe can be explained in terms of differences between your student groups rather than in terms of your educational innovation.

Be specific in your choice of expected or desired outcome, as this will help you to decide what you are wanting to measure.

Be conscious too that measurements can intrude, distorting the very outcome in which you are interested. While you wish your outcome measures to be rigorous, they should not be so detailed and extensive that they become a burden on your participants.



high



low to moderate



low to moderate



low



low

The more spread there is in the scores within the groups, the more the groups must be separated in order that we be convinced that the differences we see are unlikely to have been the result of random fluctuations in our measurements. In short, the more confident we can be that the difference we see is "real".

5. Analyse your data.

Analysis of data gathered through an experimental approach will most likely focus on deciding whether your innovation has had the predicted effect. Is there a difference to be seen in the outcome measure(s) gathered between your original (control) and post-intervention (experimental) situation? Is the difference in the direction which was predicted? And is the difference greater than a change which might be expected by random chance alone; that is, is it statistically significant?

Do not think about statistical significance as being an all or nothing thing but as an expression of your confidence in coming to a particular conclusion or making a particular claim.

Always begin the analysis with a general exploration of your data. Consider using confidence intervals first, as a good general comparison between datasets. If it appears that differences do exist, then proceed to some test of statistical significance.

Descriptive statistics (like an arithmetic mean) can be calculated, or some graphical technique (such as the plotting of a histogram) can be employed to display differences between your baseline (pre-intervention) and novel (post-intervention) measurements. Inferential procedures enable the exploration of the statistical significance of such differences. Basically, these latter procedures enable you to express the size of the differences between two (or more) groups in relation to the spread of the individual measurements within the groups.

Remember that differences in average value are not the only possible interesting outcomes. Difference in the spread of scores may be equally revealing. For example, if a topic is "well taught" (whatever that might mean) could very well result in a decrease in the spread of performance across the student group. Statistical techniques exist to explore changes of this sort as well.

Other Relevant Pages

Recipes

- Pre and post testing
- Focus groups
- Confidence logs
- Interviews
- Questionnaires
- Resource questionnaires
- Ethnography
- Trials

Information Pages

- Statistics questions
- Student sample
- Isolation or integration
- Pedagogic toolkit
- Pre and post testing
- Working with groups

Evaluation Report: Outline

References

“Beware of testing too many hypotheses, the more you torture the data, the more likely they are to confess but confession obtained under duress may not be admissible in the court of scientific opinion”

Stigler (1987)

Chris Jones
Liverpool John Moores
University.

Uses

◆ "Illuminative" evaluation:

Ethnography is useful for long-term and longitudinal studies of programmes. It concentrates on accounting for the observed setting rather than comparison between settings.

◆ Evaluation of use and the user:

Ethnographic evaluation can investigate the user's point of view. It can help in forms of participative design eliciting the point of view of those who will use the system, both educators and students. In particular, ethnography can draw out 'tacit' knowledge, the taken for granted aspects of work often overlooked.

◆ "Nth phase" evaluation:

This is an iterative design process in which ethnographic research plays two roles. First, it is used to clarify requirements prior to the design of a new system. Then, it is used to provide continuous feedback for redesign and modification.

Process involved

1. Outline

The essence of an ethnographic approach is 'in situ' observation. The events being observed should be as little disturbed and as authentic as possible. Sources of information may include:

- ◆ Naturally occurring documentation.
- ◆ Participant observation.
- ◆ Audio and video recording.
- ◆ Field notes and transcribed conversation.

2. Practicalities

Analysis and Data Collection

Analysis and data collection are not distinct phases, they occur simultaneously. Both are 'messy' and involve the use of human beings as the instrument of observation. Ethnography relies upon detail to convey the feel as well as the facts of an observed setting.

A "thick description" provides context that interprets observed actions and provides meaning to these actions. In this way analysis can fuse with description.

A "thin description" would only detail events whereas a thick description would try to analyse possible intent and the interpretation of events by participants. As a result, verbatim quotations can be the most identifiable feature of an ethnography. This can give a sense of immediacy as well as providing a record for judging any interpretations made.

Access to any situation has to be negotiated. This often involves 'gate keepers'; people who can help to introduce you to the setting with little disturbance. A good introduction can be the key to success.

Be careful to be clear when you know enough. Ethnography presents a problem of completeness, in sociology and anthropology. Fieldwork can take years. For an evaluation a rule of thumb could be feeling you know what's going on having a sense of the routine and especially knowing what you don't know or haven't seen.



low



low



high



high



moderate

Be discrete. You will come to know parts of a setting unknown to others, especially those with authority. Your aim is to gather information without disturbing the setting you are in. Often this means keeping knowledge gained from one source unknown to others.

The Role of the Observer

A central feature of ethnography is the stance taken by the observer. Indifference characterises ethnography. In this context indifference expresses an uncritical stance that doesn't set out to judge what is observed. This 'natural attitude' can be essential in gaining and maintaining trust.

Reporting

Because information may be embarrassing or jeopardise individuals the researcher must be careful to make reports as anonymous as possible. This can prove difficult if the research is 'in house'.

Variations

The level of analysis can vary greatly. This might vary from a conversation analysis (fine grain technique) to a broader analysis concentrating not only on the detail of specific interactions but also on the context and general features of the setting for the study.

The technique is essentially descriptive. When used for evaluation a balance has to be drawn between neutrality and judgement.

Full ethnographies are extensive and can take years of observation and analysis. Evaluation involves a trade-off between time taken and adequacy of description. Evaluative use generally involves 'quick and dirty' techniques.

Other Relevant Pages

Recipes

- Designing experiments
- Observation techniques
- Trials
- Nominal group technique

Information Pages

- Isolation or integration
- Pedagogical toolkit
- Pre and post testing
- Statistics questions
- Student sample

Serving suggestions

References

Innocent or informed ethnography? Ethnography can begin with either a novice or an expert set of assumptions. For evaluation purposes it is better to begin with an innocent approach that doesn't assume knowledge of what is important. Evaluators do not know in advance what will turn out to be important or interesting, even if they feel they have some good ideas. The strength of ethnography often lies in the surprise finding.

Don't go native. Immersion in a field can lead to the researcher adopting the attitudes and outlook of the participants. Blending in has to be balanced against conversion.

Involving your participants:

Before the start of the evaluation, explain what your evaluation study is about and how the data you collect is going to be used. Also try to make it clear to students that you are evaluating the software and not them.

Allow some time for people to ask any questions before and after an evaluation session and make sure that they know what is to be expected of them if they become involved in your study. If you don't know the interviewees and they don't know each other, you might like to have coffee or tea first so that you aren't going into the evaluation phase of the session as complete strangers to each other. You might also feel that it is necessary to allow some time when the students can be in the room without you.

Before you start asking questions relating to the evaluation, try and relax the participants so that they feel comfortable both talking to you and the rest of the group.

Try and value people's opinions and thank them for taking the time to become involved in your study. If possible, feedback any data collected and act on the relevant comments made during the evaluation study.

The venue for your evaluations sessions:

Try and organise an appropriate venue for your evaluation study. If you are planning to have a discussion session, select somewhere where students are going to feel comfortable and able to discuss their feelings: a shared computer lab might be the only place where you can run your practical session, but a seminar room with access to the software on a laptop, is going to be far more conducive to an evaluation discussion afterwards. When carrying out a pilot evaluation study, you might also want to select a similar environment to that which you are planning to use with the full class. This will enable you to evaluate the practicalities of your planned implementation: can all the students access and use the software? is there space for groups of students to work together?

Timing your evaluation study:

It is important to try and plan your evaluation study as far ahead as practical. This allows you to take course timetabling, reading weeks, examination weeks and holidays into account. It can be annoying to find that group of students have disappeared off on holiday when you just found time to be able to carry out the second part of your study. In addition, students are going to be less inclined to turn up at the very beginning/end of terms or just before exams.

The number of times you involve students in an evaluative study of any sort can also influence your findings: students, particularly first year students can get questionnaire fatigue. Asking students to complete a questionnaire after a series of questionnaires can result in a lack of interest or quality in their responses, particularly if there wasn't any follow-up action to their previous recommendations.

The timing of your evaluative study relative to your teaching intervention can also affect your findings: too early and your students might not have the appropriate knowledge, too late and they might have forgotten how they felt while using a piece of software. Too often and you might miss any gradual changes, too few times and you could miss a sudden change.

Issues of personality:

Carrying out evaluations necessitates a level of trust e.g. between you and your students, you and the lecturer whose pet project you are evaluating or you and someone else's students. Some students seem quite happy to express their ideas and viewpoints whereas others can appear wary of criticising or pointing out problems. The way in which you relate to the individuals involved in the study can also influence the quality of the data obtained. In addition, the students' perception of the course in which the technology is going to be used or even how they feel about the lecturer involved in teaching this part of the course can colour the way in which they feel about the evaluation session.

Rewarding your students:

Some projects pay their students to take part in their evaluation studies. For many departments, this is likely to cause various administrative if not practical problems. If it is felt that some reward is necessary for your students, you might want to supply tea and coffee before and after the session or to even organise some cheese and wine for afterwards. Alternatively, you could consider approaching publishers to sponsor the purchase of course textbooks to give your students or you could ask the university to provide them with free car parking permits, if appropriate. Generally, you will find that students are willing to participate unrewarded particularly if they feel that there will be some follow up action taken.

Named or anonymous participants?

Various arguments can be put forward as to whether you should ask students to give their names when for example asking them to complete questionnaires. Some students might lack confidence or not feel comfortable voicing their opinions to an evaluator whereas others might only take the responsibility of taking part in an evaluation seriously if they have to put their name to their opinion. If you give your students the option of putting their name on the top of a questionnaire, then generally a large proportion will leave that part blank. This puts obvious limits on a follow-up study of individuals but it is still possible to monitor any class trends.

Sometimes students are more comfortable citing their matriculation number than giving their name. Another way of identifying students is to ask them to use a personal password, unfortunately, this can result in students either forgetting or using a different word in subsequent sessions. Alternatively you can ask students to create their own password for use in your study. This can be done by, for example, using their mothers initials followed by the number of brothers/sisters they have, then their house number etc. This method works as long as the combination used will provide a series of digits exclusive to each student. Students can be reminded about the construction of their personal password each time they are completing a questionnaire for you. This also means that you can follow-up students at different stages of a course and that the students can still maintain a feeling of anonymity.

Jen Harvey
Implementation Support Consultant,
LTDI, Heriot-Watt University.

Erica McAteer
University of Glasgow

Focus groups are moderated meetings of 'involved' people discussing their experience of an educational intervention. They are a useful tool for formative/developmental or summative/retrospective evaluation and can serve as a single, self-contained method or link to other evaluation activities.

Uses

Main advantage: obtaining a large amount of interactive information on a topic comparatively easily, within a short time.

- ◆ Generating hypotheses
- ◆ Developing interview schedules
- ◆ Identifying key issues
- ◆ Developing emergent themes
- ◆ Illuminating quantitative responses
- ◆ 'Learner centred' course development
- ◆ Getting reflective feedback on interim interpretations of study outcomes

Main disadvantage: the setting is not 'natural' but deliberate.

Process

1. Define issues for focus

Start with broad themes which can be made explicit to the groups, keep a checklist of individual points of concern to prompt for, if they don't emerge naturally.

Participants will (usually!) be volunteers, it is a good idea to provide lunch!

2. Identify participants from relevant population

Try to make these representative of various types of 'user', i.e. different motivations, different entry levels, different learning environments ...

As moderator, bring in an appropriate degree of 'orchestration' – too much direction strangles disclosure and discussion. Too little, and leaders will emerge from the participant group.

3. Design the sessions

Set group size – between 6 and 12 is recommended.

Decide whether mixed groups or contrastive groups will best serve your need, comparing similar groups to check agreement or distinct groups to establish perspectives on issues.

Decide on structuring strategy – one or two broad topics, or a guided programme for discussion within allocated timeslots? Let conversation flow, if purpose is exploration. Eliciting sequenced 'rounds' of input if there is already a strong agenda to the study, if extremes within, align according to need.

The data will be 'rich' so it is best not to set too many focus items for one sitting.

Define required analysis level – qualitative, 'ethnographic' or systematic content coding, or a combination of these, depending on goals and resourcing.

Decide on recording options – notes? audio recorded? video-recorded?



Whether orchestrating or directing, try not to get drawn in to discussing the issues yourself. Don't play dumb, but acknowledge your need to learn from their experience, and listen!

4. Stage the sessions!

The most important thing is to be both confident, and relaxed – then they will be too.

5. Transcribe records

Verbatim, expressing as written text, or noting against pre-defined criteria whilst listening to/watching tape.

6. (Code and) analyse transcripts

Start with two, and examine closely to establish most useful breakdown of detail in terms of evaluation targets. Broad possibilities are by topic theme, or by participant type. When procedure agreed, test it against remaining transcripts – does it cover the data?

Whilst interpreting, and before reporting findings, it is important to check back with at least some participants to see if their perception of the 'end position' of a meeting concurs with yours.

7. Interpret findings

Integrate with other outcomes from other methods used.

Report your recommendations.

Variations

Multistage groups, where participants are brought back for more than one session, with comparison between first and later sessions. Good for developmental evaluation.

Second-order groups, mixing participants from different previous groups, where the structuring could relax slightly to see what surfaces from earlier sessions.

Other Relevant Pages

Recipes

- Nominal group technique
- Concept maps
- Interviews
- Confidence logs
- Planning trials

Information Pages

- Transcribing
- Interviewing
- Student sample
- Pedagogical toolkit
- Isolation or integration
- Working with groups

Serving suggestions

References

Erica McAteer
University of Glasgow.

"Interviews are conversations where the outcome is a coproduction of the interviewer and the interviewee" Kvale (1996).

Uses

There are a number of approaches to interviewing, which have in common the need to get perspectives on the evaluation targets from a sample of 'users' representing different stakeholder groups.

- ◆ The standardised, open ended interview

Strength: makes sure questions are asked in the same way across a sample population by different interviewers.

Weakness: risk losing important, unanticipated, information.

- ◆ The guided or structured interview

Strength: keeps interaction focused, covering same ground with respondent sets, while allowing individual experience to emerge.

Weakness: cannot divert far, or long, from agenda without losing part of 'the story'.

- ◆ The informal, conversational interview

Strength: allows the evaluator to respond quickly to individual differences and situational changes.

Weakness: a great deal of time is needed to get systematic information.

Like Focus Groups, interviews can be useful for formative/developmental or summative/retrospective evaluation.

Process

1. Define your population sample

Who (and how many) will you interview?

2. Identify the target interview issues

List them by 'theme' if possible.

3. Design your study

Think of how to approach interviewees, the conduct of the interview itself and allow time for post-interview feedback.

4. Produce an interview script

Broad cues for a conversational approach, guide-list for a structured interview, questions for a standardised 'open interview'.

There are several sorts of questions that can usefully be asked, and Patton (1996) suggests that these should if possible be sequenced in the following way: behaviour/experience questions; opinion/value questions; feeling questions; knowledge questions; sensory questions; demographic/background questions.

5. Pilot the interviews

Use a small subset (at least two) representatives of the different populations involved. This is an often skipped but absolutely essential step!

Asking exactly the same questions of very different groups can be very useful, particularly when a development has to meet circumstances of need so far as teacher, learner and institution is concerned – but it is not always appropriate to predict the answers will be different!

6. Conduct the interviews

Write up a short note after each episode, including anything that stood out, went amiss, etc. just logging general procedure.

Probes, follow-up questions, clarification requests can be made as and when necessary, taking care to maintain a light, interested touch or you might stem the flow.

7. Transcribe responses and organise in best way for analysis

This can mean tagging 'rich' text responses to secure context information (respondent details, intervention details, schedule details) and storing responses under issue/question headings for whatever level of content analysis is appropriate. It can mean providing for 'response profiles' from individual participants. It is quite likely that both orientations will be necessary.

Typing each interview response into a spreadsheet with the individual case as a row with columns holding demographic information, respondent factors etc. having one column for each open question is a good way of holding the data safe. Printouts of individual columns with all the text responses can then be coded and this information entered into a duplicate file. Excel 7 allows word-wrap for quite lengthy responses!

If you can, check back with interviewees to obtain elaboration, further comment, confirmation (or refutation!) of your interpretation of their meanings.

8. Analyse outcomes

Remember that this technique gives qualitative data! Even if it has been possible to code responses such that descriptive statistics are possible and useful, any further analysis has to be very carefully conducted, using 'non-parametric' tests.

9. Interpret findings

Integrate with other outcomes from other methods used.

Report your recommendations.

Variations

Phone interviews, where the interviewer fills in the responses as they come over the wire – still allows prompting and 'real-time' interaction.

Online interviews, either using Web-forms or just through email. Here some structure is essential with an introductory paragraph if agreement not already obtained before sending. What is lost in 'immediacy' might be balanced by 'reflection', with a two-shot message circulating feedback and seeking further comments.

Tape interviews (Lockwood) – sometimes useful to use different methods with exactly the same questions as the responses do differ for different modes.

Self-recorded audio taped interviews, where the respondent is sent a list of questions or issues and asked to record responses at a suitable time (perhaps at the end of a course of study) is an unusual, but workable, alternative.

The above, and the obvious corollary of the 'paper interview' shows that the line between interviews and open questionnaires blurs at the edges...

Other Relevant Pages

Recipes

- Questionnaires
- Checklists
- Designing, experiments
- Trials
- Focus groups

Information Pages

- Transcribing
- Interviewing
- Student sample
- Pedagogic toolkit
- Isolation or integration
- Working with groups

Serving suggestions

References

When you are carrying out interviews, focus groups and discussions it is important to keep a record of what was said for future transcription, reference or analysis. The method you use will depend on how much detail you require.

Keeping written notes:

Taking notes as you go along can be a useful way of gauging general opinion and the time taken to stop and take notes can be a useful way of giving breathing space during the discussion. However, trying to keep the discussion going and keep an accurate record of everything said can prove to be an almost impossible task.

Taping your session:

The easiest way to keep a record of a discussion, is to tape your session. However, it is important to make sure that you have your participants' permission to tape their interview. In addition, let them know exactly why you are carrying out this exercise and what you are planning to do with the information. You might also like to let the participants see any transcriptions that you are planning to use.

Sometimes students can feel awkward about the idea of being taped so you might start by having some general conversation at the beginning of your session in order to get them accustomed to talking with a tape recorder running.

Hiding a tape-recorder and recording secretly might seem to be a good tactic to help keep your participants relaxed, but apart from the obvious ethical problems this raises, it also means that legitimately, you can't make use of any of information you collect.

Setting up the taping equipment:

Make sure that you are familiar with using the equipment and that everything is working OK before you start. It can be very disruptive to a session, if you have to spend the first 15 minutes trying to get the tape recorder to work. Batteries in the tape recorder or tapes running out can also mean that you lose important sections of a discussion.

Having a good microphone is probably more important than having a good tape recorder. Constant hissing on a tape or just having indistinct mumbling can cause all sorts of problems when transcribing. Placing a microphone on a soft surface such as foam or a piece of carpet can help to improve sound quality and a flat microphone can also appear less intrusive.

Most microphones built into recorders have a restricted range and if you are working with a group of people, you could find that you miss one or two people's contributions because they are just out of range. If you are working with a group, try to make sure that they are an equal distance from the microphone and if you can't get hold of an omnidirectional microphone try and use a second one.

Recording progress:

It is often a good idea to try and give your groups a feel of their progress during the session. This can be done through using a spray diagram or flip chart to record ideas as you go along. Visual methods can be more effective than a list of points. They also have the benefit of focusing discussion and ensuring you are accurately interpreting what is being said as any misconceptions can be corrected there and then. In addition, when an idea is represented in this way it becomes separated from the individual who put it forward and therefore it becomes easier for the other group members to criticise or disagree with its content. As part of your exercise, you could also ask your groups to summarise the main points of the discussion.

Working with recordings:

It is a good idea to try and transcribe your own tapes as you can remember more clearly who said what. A speed play option on your recorder can assist when scanning through material for particular sections and foot pedal attachments are available, whereby you can control your tape recorder being switched on and off. These are particularly useful when you are transcribing data.

If you don't want to go through the process above, you might like to use software (e.g. CODE-A-TEXT) which is now available and which works from digitised sound files. The coding, memoing and analysis can be done in the software while listening to the sound file. This software can also work for transcribed data primarily as a content analysis program.

Identifying who said what:

Which ever method you are using it is a good idea to ask someone to keep a note of the first few words said by each participant. Reference each participant by their initials or give them each a reference number. If you don't have a scribe to do this, then think about mentioning the participant's name in the conversation after or before they speak, to make it easier for a transcription. It is important to try and do this in as natural a way as possible and in a way which doesn't affect the flow of the conversation.

Focus groups can sometimes comprise students/ staff from different disciplines which can make it more important to identify who says what. You might consider using speakers identifiers which look something like this : BSC-1-8, MA-3-6, HND-1-2 etc. representing a first year student from focus group 8, a third year student from group 6 and first year HND student from group 2. This also means that you can carry out automated text searches on the basis of the speaker's identifier if you are using data analysis software at a later stage.

Jen Harvey

Implementation Support Consultant,
LTDI, Heriot-Watt University.

Gaye Manwaring
Northern College.

An interesting method which can yield surprising results. Advantages are that the views of everyone are considered, and the time for all stages is low. It is a group in name only as the views of all participants are collected separately.

Uses

- ◆ Identifying key issues to be explored by other evaluation methods
- ◆ Identifying issues that concern many people
- ◆ Useful for formative evaluation or review
- ◆ Good at giving participants a sense of ownership

Process

This must be done without discussion so each individual's views are included.

1. Individual thought

Ask each person to note down their own views or reflections in response to the statement or question from the evaluator eg "List three best and three worst aspects of the course"; "What is the most important aspect to be changed?"

Allow only a few minutes.

If the situation is delicate, the evaluator could leave the room for this stage, so no-one is identified with their comments.

2. Collect responses

Go round the group and ask everyone for a comment. List these on a flip chart.

Take everyone's first answer before allowing anyone a second choice.

Cluster similar items together but do not discuss issues.

Let people write their own votes on the flipchart so their scores are anonymous.

3. Vote

Allow each person to allocate six votes to cover their views. They can give all six votes to the item they feel most strongly about or they could give one point to each of six items or four points to one item and two to another, etc. This identifies items with high scores and those with low scores from many people.

4. Discuss (optional)

Put people into groups to discuss the most significant items – analyse problems, propose actions etc.

Give each group a different item. Or give the same item to a group of students and a group of tutors.

Variations

Snowballs

After stage 1 put people into pairs and get them to agree to a compromise list. Then put the pairs into groups of four and again ask for discussion and a consensus. If there is time, continue until there is just one group. If you ask for, say, three items, make sure the list does not get longer when you move to larger groups.

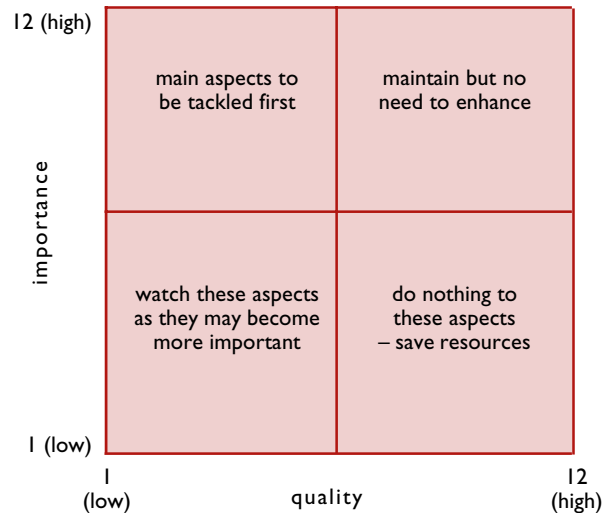
Prepare the grids in advance.

Quality/ importance grids

In stage 1 ask for the identification of key aspects of the course. List and cluster on the flipchart but you must end up with exactly 12 aspects. Then each person is asked to rate each aspect from 1 - 12 for quality and importance and then to put these on a graph.

Then ask groups to discuss the aspects in the top left quadrant.

Ensure that the origin is one, not zero.



Other Relevant Pages

Recipes

- Focus groups
- Interviews
- Concept maps
- Trials
- Cost effectiveness

Information Pages

- Working with groups
- Pedagogic toolkit
- Isolation or integration
- Student sample

Serving suggestions

References

The idea of pre and post testing of students is often accepted as a viable method to assess the extent to which an educational intervention has had an impact on student 'learning'. However, despite the fact that the results of such evaluative measures are extensively reported in the literature (generally as part of a more extensive discussion of particular CAL interventions) there is a dearth of detail on how the tests were conducted and of the criteria which have been used to measure their impact. Furthermore, there is rarely any discussion of the potential effects of other factors as contributing to what is often reported as 'improved student performance', least of all on the possibility that the testing process itself may have a significant effect on the outcome of the evaluation.

As with all evaluations, the starting point must be to define clearly why the evaluation is being performed, what is being evaluated, when the evaluation will be performed, and how it will be performed.

Why?

Because we know that students with different skills and backgrounds come to study a particular subject, we need to establish a base measure of their knowledge and understanding of a topic in order to be able to quantify the extent of any changes in this knowledge or understanding by the end of a particular period of learning. Ideally, we wish to know not only that the educational intervention has had an impact, on the student, hopefully a positive one, but we also want to be able to quantify that impact.

What?

Simplistically viewed, the process should entail students undertaking a test to determine some identifiable starting level of knowledge or understanding of a topic and a later point undertaking an exactly comparable test to determine the extent to which knowledge and understanding has been augmented by the educational intervention. However, unless we content ourselves with assessment instruments which simply seek to measure the ability to retain and recall 'known facts', it is difficult to provide an accurate scalar measure of the exact extent of the 'improvement' in performance. Furthermore, these measures do not directly inform us of the process by which the improvement was effected. There are a large number of variables introduced by factors which are intrinsic to the measurement of performance and also those which are related to the circumstances of the subjects being assessed. When combined, these create a level of tolerance which makes significance testing using statistical analysis of the empirical data virtually impossible.

Thus in the area of pre and post testing of students it is important to realise that one must be prepared to make only very limited claims for any findings which seek to provide an accurate measure of the change in learning which can be attributed to the intervention.

When?

Also it is important to bear in mind that the timing of the pre and post tests will have a critical impact on the results obtained. In some areas of study pre testing of students is not appropriate because at the outset of the course it is already known that students would be expected to have virtually no knowledge or experience of the subject. Pre testing would be pointless from the students point of view and so the pre test really has to be done at some point in time when we can expect the student to have acquired some relevant knowledge but before the student is exposed to the CAL materials. The results of post testing will vary considerably depending on how soon the test is administered after the student has used the CAL materials and ideally if the test is administered immediately after the learning session additional follow up tests at later dates should be used to provide some evidence of application and impact of what has been learned.

How?

The design of the pre and post questions is critical to success. The repetition of the same test questions is obviously not a sound solution to achieving comparability but it is a good idea to retain a proportion of the original test materials and to blend this with new questions which examine the same expected learning outcomes. It is also important to consider the type of questions used. Certainly we should not rely purely on objective questions but, extended questions which seek to test a whole range of issues are also inappropriate. However, the use of short definitions can make it easy to identify an accurate and unambiguous response. It is also possible to consider incorporating a requirement for students to indicate the level of confidence in their response. Design of pre and post tests is not easy and above all we have to be extremely careful in the test design to ensure that we are not simply 'testing the test'.

This all seems to paint a very negative picture of the value of pre and post testing of students. However, if one is willing to accept the fact that pre and post testing should be seen as assisting us to learn more about how students use CAL rather than as a means of demonstrating changes in knowledge or skills then we can set about using appropriately designed pre and post tests to achieve this objective. What we want to generate is not simply a performance measure of the number of 'correct' responses provided by a student but the manner in which the deliver mode of instruction has caused an alteration in their responses.

Robert Newton
The Robert Gordon University,
Aberdeen

Robert Newton
The Robert Gordon University

The following recipe describes a test methodology currently being used to test a group of postgraduate students undertaking a programme of instruction in bibliographic classification, but which has broad application possibilities (see the overview in the information page).

Use

- ◆ Evaluation of 'established' CAL packages i.e. those in which issues related to content and interface design have already been tested and proven to be acceptable.

Process

1. Decide on the composition of the test group

Ideally, to perform pre and post testing, you should aim to achieve a student test group of at least 30 students. You will find that should you try to use this technique on very large student populations the work involved in determination of student profiles and analysis of test results becomes unmanageable (unless you are putting considerable staffing resources into performing the survey).

2. Create your student profiles

When creating your student profiles, try to take into account the various external factors which can influence student learning. Profiles should therefore be created to elicit detail of:

- ◆ motivational factors,
- ◆ personal factors (derived from interview),
- ◆ educational factors (including previous experience and use of computers and CAL),
- ◆ learning style.

3. Familiarisation with CAL package

Although a package might be simple to use it is important to ensure that students are familiar with all aspects of how to use the various features – such as taking notes online, activating the glossary, navigating through various pages and beginning an online test session.

You could consider organising a familiarisation session prior to your evaluation.

As part of the evaluation session, students could be asked to complete a questionnaire which includes questions relating to their degree of comfort and confidence in being able to use the CAL package.

If you intend that pre and post testing of students should be complemented by other evaluation techniques e.g. paired testing of students, you should aim to use a test group of 60 students in order to allow you to make reliable statistical conclusions from your data.

Pre and post testing is extremely sensitive to the external factors which influence student learning. It would be very wrong to assume that we can discount a variety of external factors which will affect student performance when considering the reasons for variation between pre and post test responses.

Allow time to make sure that all are comfortable using the CAL package.

Have a question at this point which asks the student if he/she would be prepared to use CAL on occasion as a substitute for lectures.

In order that students have a basic foundation of the prerequisite subject content, run the CAL session after the delivery of a lecture on an appropriate topic as one part of their formal programme of study.

4. Pre and post testing

A. work around the CAL package

Decide when you want to pre and post test your students. Think about how much of the subject content they need to know before a pre test. It is advisable to post test immediately after they have completed their study of the material in the CAL package.

OR

B selection of groups for two alternative modes of learning

One option might be to use the data derived from student profiles and responses to stage 3 questions so that the cohort can be divided into two. One group can use the CAL package as a substitute for lectures (on at least 2 occasions). The second group can follow the standard lecture programme. Both groups should undertake pre and post tests.

OR

C work around the lecture

This is the reverse of Stage A. At this stage all students take the CAL unit prior to the delivery of the lecture in the topic. The pre and post testing is delivered immediately prior to and immediately after the lecture. These tests could be online or paper-based.

Ideally there should be some form of tracking and timing mechanism built into the package being used and this will be very valuable when attempting when analysing results of tests.

5. Analysis of results

The various tests will provide a huge amount of data – some of it will be raw numeric data that can be analysed using standard statistical tests.

The advantage of having taken such an exhaustive approach to examining all the factors which might affect student learning (when building up the student profiles) is that, as was said earlier there are a large number of variables which can potentially affect student learning. At the analysis stage you will be able to take account of all of these variables.

Other Relevant Pages

Recipes

- Questionnaires
- Checklists
- Confidence logs
- Interviews
- Designing experiments
- Cost effectiveness
- System log data
- Observation techniques

Information pages

- Learning styles
- Pre and post testing
- Statistics questions
- Student sample
- Pedagogic toolkit
- Working with groups

Serving suggestions

Much of this seems to be very complex and intensive of staff time and effort. Pre and post testing of students is not an easy evaluation strategy, but it does provide a great deal of useful information on how students learn.

Kirsty Davidson &
Judy Goldfinch
Napier University.

Uses

- ◆ To obtain feedback on some activity
- ◆ To obtain views and/or factual information from people

Process

Think whether a questionnaire is the best way of getting the information or whether you can use direct observation or a carefully controlled experiment.

1. Define your sample

State your objectives clearly.

Decide whether data can be collected on everyone involved (a census) or whether you are restricted to collecting information from a smaller group (a survey).

For surveys:

- ◆ define clearly your population of interest;
- ◆ think carefully how you can obtain a 'representative' sample (it is very easy to introduce bias into a survey if the sample is not chosen carefully) and how big your sample should be (the size will affect the reliability of the results).

2. Clarify your objectives

Be very clear about your objectives and this will make it easier to write down precisely what information you require.

3. Develop your questionnaire

a. Construct statements

Think about whether you want to use open-ended questions or closed questions or a mixture of the two. Open-ended questions allow the respondent to express an opinion on some matter and could be a single word or a long statement; closed questions require a specific answer e.g. a box to be ticked, items to be ranked, etc.

Likert-type scales are useful for assessing a respondent's attitude to a statement, e.g. strongly agree, agree, neutral, disagree, strongly disagree. In order to avoid respondents agreeing with a set of statements without thinking clearly about each one, vary positive statements (e.g. I found the package easy to use) with negative statements (e.g. the screens were too cluttered).

b. Design the questionnaire:

- ◆ Include an introduction explaining the purpose of the questionnaire
- ◆ Give clear instructions on how to fill it in
- ◆ Keep questions simple and unambiguous
- ◆ Use a consistent style (not a mixture of ticking boxes, circling answers, etc.)
- ◆ Arrange questions in a logical order putting sensitive issues towards the end
- ◆ Include a thank you at the end
- ◆ Including codes for responses can be useful if the analysis is to be carried out using a computer package

Responses to open-ended questions can be very useful in preliminary studies to find out what issues respondents consider to be important; however they are much more difficult to code and analyse.

If possible, make use of or adapt a standard questionnaire which has been tried and tested by other researchers.

Keep the questionnaire as short as possible.

Try and give the questionnaire a professional appearance.

Allow time at the end of a class to get students to complete a questionnaire while they are captive!

4. Practicalities of application

Before using the questionnaire in earnest, pilot it with 4 or 5 respondents to uncover any potential problems in the questionnaire and in the analysis.

Questionnaires can be:

- ◆ handed out and completed there and then (ensuring a good response);
- ◆ handed out and collected in at some later date (giving respondents more time to consider their answers);
- ◆ posted out (generally results in a poor response rate);
- ◆ administered in person or by 'phone (takes much longer but achieves a good response rate and allows clarification).

Consider whether non-respondents would be likely to answer the questions differently to the respondents perhaps causing bias. Follow up non-respondents if possible.

5. Analysis

Check answers for values outwith expected ranges and inconsistencies.

Check the answers to open-ended questions to see if they have common responses that can be coded up in some way.

There are many ways to usefully present the information from questionnaires visually – bar charts, histograms, scatterplots, etc. Most statistical analysis packages allow these to be produced easily. Spreadsheets can also be used to analyse questionnaire data. However, if a very large survey is being undertaken, it may be worth considering a specialist package such as SphinxSurvey, distributed by Sage Publications (<http://www.sagepub.co.uk>).

Example from a questionnaire

Finally, please give us your views on using CBL packages to learn mathematics

a) What drawbacks do you see in using CBL?

b) Can you suggest ways of overcoming these drawbacks?

c) What benefits do you see for using CBL?

Other Relevant Pages

Recipes

- Checklists
- Pre and post testing
- Trials
- Designing experiments
- Resource questionnaires
- Interviews
- Focus groups
- Confidence logs
- Cost effectiveness

Information Pages

- Likert scales
- Questionnaires
- Guidelines for questions
- Statistics questions
- Student sample
- Interviewing

Serving suggestions

References

All lecturers come in contact with questionnaires whether it is the standard end of year course questionnaire or one that is used in research. These questionnaires come in many different forms from: factual to opinion based, from tick boxes to free text responses. Whatever their form, questionnaires are often viewed as quick and easy to do. This is not always the case. To get useful responses, in a cost-effective way, it is important to be clear about the aim of the questionnaire and how the responses will help you improve the learning technology or its implementation. Think also about the analysis of results. It can be sobering to consider the amount of data you will generate and the time it will take to analyse.

Some advantages and disadvantages of questionnaires follow. Notes on how to deal with some of the disadvantages are also provided, as are some references to more comprehensive information on questionnaires.

Some disadvantages of questionnaires:

- ◆ Questionnaires, like many evaluation methods occur after the event, so participants may forget important issues.
- ◆ Questionnaires are standardised so it is not possible to explain any points in the questions that participants might misinterpret. This could be partially solved by piloting the questions on a small group of students or at least friends and colleagues. It is advisable to do this anyway.
- ◆ Open-ended questions can generate large amounts of data that can take a long time to process and analyse. One way of limiting this would be to limit the space available to students so their responses are concise or to sample the students and survey only a portion of them.
- ◆ Respondents may answer superficially especially if the questionnaire takes a long time to complete. The common mistake of asking too many questions should be avoided.
- ◆ Students may not be willing to answer the questions. They might not wish to reveal the information or they might think that they will not benefit from responding perhaps even be penalised by giving their real opinion. Students should be told why the information is being collected and how the results will be beneficial. They should be asked to reply honestly and told that if their response is negative this is just as useful as a more positive opinion. If possible the questionnaire should be anonymous.

Some advantages of questionnaires:

- ◆ The responses are gathered in a standardised way, so questionnaires are more objective, certainly more so than interviews.
- ◆ Generally it is relatively quick to collect information using a questionnaire. However in some situations they can take a long time not only to design but also to apply and analyse (see disadvantages for more information).
- ◆ Potentially information can be collected from a large portion of a group. This potential is not often realised, as returns from questionnaires are usually low. However return rates can be dramatically improved if the questionnaire is delivered and responded to in class time.

Further Reading

There has been a lot written about questionnaires. Listed below are some sources of information that you may find useful.

Carter, MP and Williamson, D (1996) Questionnaire Design. Staffordshire University Business School, Leek Road, Stoke-on-Trent ST4 2DF, United Kingdom http://www.staffs.ac.uk/buss/bscal/mandev/m_qm/t_que/que.htm

Gilbert, N (1993) Researching Social Life. Sage Publications, London.

Kirakowski, J (1997) Questionnaires in Usability Engineering. A List of Frequently Asked Questions, Human Factors Research Group, Cork, Ireland. <http://www.ucc.ie/hfrg/resources/qfaq1.html>

Moser, CA and Kalton, G (1979) Survey Methods in Social Investigation. Gower Publishing Company, Aldershot, England.

Oppenheim, AN (1992) Questionnaire design, interviewing and attitude measurement. Pinter, London.

Wilson, N and McClean, S (1994) Questionnaire Design: A Practical Introduction. University of Ulster. Copies available from: UCoSDA, Level Six, University House, University of Sheffield, Sheffield S10 2TN.

Tel: 0114 272 5248, Email: j.burgan@sheffield.ac.uk

John Milne
Centre for CBL in Land Use
and Environmental Sciences,
Aberdeen University.

General advice

Try to ask questions only directly related to what you are evaluating and not just for the sake of it. A few focused questions are much more useful than a collection of general ones.

Make sure that the student knows how you expect them to respond to a question e.g. do they have to tick/cross underline/write their own answer etc.

Avoid double-barrelled questions, as students might want to respond differently to each part e.g. 'How did you feel about using packages A and B?'

Try to keep the language simple and easy to understand as students will often miss out a question which they don't understand.

Make sure that your questions aren't ambiguous and open to different personal interpretations e.g. a question like 'Do you prefer tutorials or simulations?' is dependant on what a student perceives as being a 'tutorial' and a 'simulation'. Including a definition of each would increase the question's clarity.

Try to keep the questions short, as long questions can sometimes be confusing.

Try to provide guidance in the length of answer you would like the student to give and how long it might take them to complete all the questions.

Avoid questions with obvious answers or ones which are likely to elicit the same answer from all students or ones fishing for compliments e.g. 'How do you rate this course?' 'brilliant', 'marvellous' or 'best ever'.

Start by asking your more straightforward questions and then work through to those requiring longer answers.

Vary the question formats in order to encourage students to think about each of their responses and not just go through ticking 'agree' boxes.

Avoid yes/no questions unless you want a yes/no answer.

Avoid 'Why?' questions or strongly evocative questions which might make a student defensive.

Group questions investigating similar themes together, perhaps using a header.

Fixed response questions

Try to balance your order and use of negative and positive statements.

Make sure that your response options are mutually exclusive and don't overlap.

Try not to use negative wording in question statements as this can often lead to double negatives when added to several response options.

Try to put your responses into a logical order if one exists – this enables a student to make more of a relative judgement.

Try to make sure you include the full range of possible answers. The use of 'Other' and 'Please specify' as an option gives a student an opportunity to add in their own response.

Watch if you use 'don't know' as a description for the midpoint to your scale. This could indicate either that they don't understand the question or that they just don't want to state an opinion. To avoid misinterpretations you can always suggest at the top of the page that students miss out any questions they don't understand.

Watch the phrasing of responses given in the form of an attitudinal scale. It can often be difficult to provide a well balanced range of descriptive words in an attitudinal scale. A Likert scale using numbers or points on a scale between phrases like 'strongly agree' and 'strongly disagree' can give a better range of options. Alternatively, instead of limiting their options, you might ask the students to indicate their response by placing a cross on a line in the appropriate position between opposite opinions.

Providing a Likert scale with an even number of options can encourage students to make a decision, but sometimes can result in students just adding in a midpoint option themselves.

Jen Harvey

LTDI Implementation Consultant,
Heriot-Watt University.

Robin Shaw
University of Glasgow.

Uses

- ◆ Find out what resources the students are actually using
- ◆ Determine how much time they spend on each resource
- ◆ Establish how much value they place on them and any difficulties they experience in accessing the resource

Rationale

However good a CAL package is, it will have no positive effect on student learning if it isn't used. Students may not have been able to access a CAL package due to a shortage of computers or computer lab. hours, they may lack the time or the motivation to attempt to use it, or having used it they may not have found it valuable. Effective integration of CAL or any other innovation into a course is essential otherwise no matter how good it is, it will not be of value to students. The questionnaire can be useful in providing information about resources other than the CAL package whose use it was designed to illuminate.

Process

1. Identify resources

Find out from the lecturer what resources are available to the students.

Have a focused talk with the students about how they collect information.

2. Decide what information you want to gather about the resources

This might be, for example: resources used, their usefulness, the number of times they were accessed, the difficulty of access, etc.

Limit yourself to the most important resources to prevent creating overlong questionnaires.

3. Construct and administer the questionnaire

Remember to allow enough time for thought and accuracy by students.

Example of a Resource Questionnaire

Resource	tick if used	not at all useful	not very useful	useful	very useful	extremely useful	Reason for answer
lectures	✓		✓				
tutorials	✓				✓		Best part of course

4. Analyse the results

The results can best be collated by coding the responses on a spreadsheet either using the numeral 1 where the student has ticked the questionnaire or by using the numbers 1 (not useful at all) to 5 (extremely useful) to indicate the response.

Spreadsheet Coding 1

Resource	tick if used	not at all useful	not very useful	useful	very useful	extremely useful	Comments
lectures	✓		✓				
tutorials	✓				✓		Best part of course

Spreadsheet Coding 2

Resource	Most frequent response code	Comments
lectures	2 [the other responses can be added in this same column, and analysed by producing a frequency count for each possible response.]	
tutorials	4	Best part of course

Presentation of the results

When all the responses have been entered from the questionnaires, the results can be summed and the comments printed separately along with the resource and response which inspired them. Then the information can be presented in tabular form, as a bar chart or as a pie chart. Alternatively, the results can be displayed on a copy of the questionnaire.

From an analysis of the information, action can be taken to ensure that valuable resources are properly used. For example, the questionnaire may reveal the need for the timetabling of computer access, or perhaps for a clearer introduction to the package and how it is to be used. It will frequently show that where the package is not fully integrated in the course but is seen as an option, it will be used lightly. This should not surprise us; students like lecturers are busy people and seek strategies which have worked in the past to pass the exam.

One of the valuable aspects of the questionnaire is that it covers the entire suite of resources available to the student and therefore will also highlight problems with traditional resources as well as with innovations. So it is a worthwhile exercise even within a course where no new resource has been offered.

Other Relevant Pages

Recipes

- Questionnaires
- Checklists
- Pre and post testing
- Interviews
- Trials
- Designing experiments

Information Pages

- Pedagogic toolkit
- Guidelines for questions
- Likert scales
- Statistics questions
- Student sample
- Working with groups

References

Kirsty Davidson
Napier University.

Uses

- ◆ Investigating a piece of software under development
- ◆ Introducing students to the use of learning technology
- ◆ Investigating students' usage of computer-based materials

Process

Try and time the session so that the student has covered at least some of the material in the software.

1. Identify your situation and/or materials for investigation

These could be:

- ◆ learning technology materials still under development into which formative evaluation could feed;
- ◆ piloting software in order to provide information for a subsequent implementation strategy.

Check availability of studio, computers with the software and technical help. Pressgang two to four student volunteers.

2. Identify key areas of interest

These can be:

- ◆ the usability of the software;
- ◆ the student's perception of the software as a learning tool.

Ensure the task is achievable within about half an hour.

3. Plan your evaluation session

Decide whether you or another person, acting as a facilitator, are to be running the session.

Try to select tasks which explore different features of the software and which are relevant to its anticipated usage with students.

Assure the students that it is the software which you are evaluating, not them. Sit beside them and encourage them to speak out their thoughts with prompts like 'why did you do that then?', 'what are you thinking just now', 'what do you think of...!', etc.

4. Practicalities of application

Check the timing of the session against the students program – volunteers are not likely to be so forthcoming at the start of term or the day before an exam.

Check the student knows where the studio is and can get there and back in time for classes.

Before the session check the software runs OK on the computer in the studio.

Two cameras should be set up, one pointing at the students upper body to catch face and hand movements, the other pointing at the monitor. These images are then combined into a split screen video.

A microphone should be attached to the student and possibly the facilitator so that sound is captured.

Ask the technicians to overlay a timer on the video to make it easier to locate specific sections etc.

5. Analysis

At least a rough transcription of the video is useful for reference purposes.


The video can be viewed as often as required for different purposes.

Snippets are very useful for demonstrating a point during a talk and/or livening up a conference presentation

Variations

This technique could be used for evaluation of, or staff development in, any form of independent learning.

Example screen shot from a split screen video



Other Relevant Pages

Recipes

- Supplemental Observation
- Interviews
- System Log Data
- Ethnography
- Trials
- Designing experiments

Information Pages

- Learning styles
- Transcribing
- General Issues
- Interviewing
- Working with groups
- Student sample

Serving suggestions

References

The manner of transcription may matter whatever techniques are being used to analyse textual data, but consistent transcription, following a basic set of rules will be even more necessary when using a CAQDAS (computer assisted qualitative data analysis software) program, or even word processing search tools, to assist in the analysis process.

SECTION 1. Aims to list some basic Do's and Don'ts – which will help reduce the necessity for later changes to data or preparation time – whatever software package is being used. It would be helpful to give this list to a transcriber – even if you are not certain that a CAQDAS software package will be used.

SECTION 2. Lists particular points which are software specific, without placing too much emphasis the full extent of instructions which might be applicable e.g. for data preparation using NUD*IST (1). This aims to be a short list providing pointers only. More detailed information will be available in manuals or from software specific workshop curricula.

1. General Guidelines

These are of especial relevance when considering the use of software to assist in the analysis.

Data are often inconsistently transcribed within one file. Much of the final preparation which may have more to do with analysis stage decisions and will have to be left to the researcher. However he/she will be saved a lot of tidying up work if some simple basic minima rules are applied at an early transcription stage. This list is general and NOT 'software-specific'. So these are not all the steps required for the preparation of data for individual software packages.

Some of the guidelines below are even more relevant for data with structure – e.g. surveys containing open ended questions, focus group data, structured interviews, etc.

NOTE: none of this need apply to data being prepared for Code-a-Text (2), which can use rich text, or word processed files.

DO – always make spelling, spacing etc., of repeating speaker identifiers, question headers, section headers, topic headers, absolutely uniform throughout text, e.g. QUI: OR QI.; NOT a mixture of both. You may need to depend on this uniformity when performing text searches and saving the results. It is easier to use text search tools which look for exact strings of characters, not approximations.

DO – create clear line space between paragraphs. Not always a requirement but will help with e.g. Atlas/ti (3).

... BUT! ...

DO NOT – have a line space between speaker identifiers, topic headers etc., and the text associated with them. (This will be more relevant in some CAQDAS software than others).

DO NOT – depend on Word processing display characteristics to indicate things (emphasis etc.) about your text e.g. bold, italic, underline, e.g. do not use emboldening or italics to identify the respondents speech.

DO – use a clear speaker identifier instead, preferably in UPPER CASE. This will allow CASE SENSITIVE searches for the speaker identifiers. E.g. IV: or CAROL:

DO NOT – use bullet points (unless you want them to act as section headers in NUD*IST – and you must understand the significance of 'Section headers' as they are used in NUD*IST before deciding this).

DO – be clear about the amount of 'preparation stage' thought necessary for your data; this may vary enormously depending on the CAQDAS software being used. Try to find out how this varies, by asking or experimenting or reading the manuals!

DO – with the above in mind, if possible before transcribing too much data, prepare a small pilot project with one or two files, inside the CAQDAS software. Do some coding and retrieval, and text searches, to test that the formatting of the data seems to work or if you see where improvements can be made – they can be applied at an early stage of transcription.

DO – save an original, 'safe' copy of data, with the above thoughts in mind, in Word Processed format. You can always go back to this if subsequent SAVE as TEXT options go wrong.

DO – transcribe in a fixed width font like Courier or Courier New. Make the points size 10 or bigger. Use the usual default margins of 3.17 cm either side or 1.25 inches. This will usually present a line length which sits well in most text windows – if not, some software packages, like Ethnograph (4) have editors, or format processes which convert the transcribed files for you into convenient line length.

2. Peculiarities of transcription requirements for a sample of individual CAQDAS software programs

ATLAS/ti textual data must be saved as 'Text with line breaks' (that is hard returns introduced automatically at the 'save as text' stage) other wise one long line stretches way beyond the available space until the first paragraph break. Additionally, do make a clear line space between paragraphs.

WINMAX Pro 97 (5) data will be better imported to the software WITHOUT line breaks. This will assist autocoding text selection later. Do not put a hard return after a Speaker identifier – begin text on the same line. Insertions of special syntax in transcript will allow the cutting of multi case files into smaller files within files. See the software manual for further details.

QSR NUD*IST will require you to make a decision about what the minimum codeable chunk/segment will be throughout your data, before data is imported into the software. The 'text unit' is defined by where the hard return goes. So using a 'Save as Text' option which introduces line breaks at the end of every line will result in LINE text units. Or interactively insert a hard return at contextual breaks or sentences, pauses etc., giving codeable segments of variable length. Or keep hard returns in all natural places, i.e. paragraphs, sections, speaker sections etc. Note: Section headers have a specific purpose in NUDIST. See the software manual for further details.

Further details about software packages

1. QSR NUD*IST: <http://www.qsr.com.au>
2. Code-a-Text: <http://www.codeatext.u-net.com>
3. ATLAS/ti: <http://www.atlasti.de/>
4. WINMAX Pro 97: <http://www.winmax.de>
5. Ethnograph: <http://www.QualisResearch.com>

All the packages are available from Scolari, Sage Publications Ltd., 6 Bonhill Street, London EC2A 4PU (downloadable demo versions available from <http://www.soc.surrey.ac.uk/caqdas>).

Ann Lewins
CAQDAS NETWORKING PROJECT,
University of Surrey.

Gayle Calverley
The University of Hull.

This involves watching how your students make use of a piece of software.

Uses: example situations

- ◆ Evaluating new resources introduced to an existing or new course or known resources within a new course structure
- ◆ Identifying student IT skill gaps for subject-based tutors
- ◆ Immediate feedback to the material developers or the course tutor is often the most useful aspect of this technique

Process

Check out the location. Check your assumptions about observation circumstances. Confirm the observation situation with the person running the session. If it is your own session then, consider recruiting a third party observer, who could be a volunteer or paid at demonstrator rates.

Would what you hope to observe be as effectively answered in group or individual interviews with the users immediately following the session? Consider directly comparing interviews with what was independently observed during the session.

Look at the interaction between the group and the person leading the session - What type of tutor - student interaction is taking place?

Make sure the proforma is designed for "at-a-glance" use, and is not too long. There may not be much time in a session for record keeping. Leave room for comments on unusual findings or situations. Checklists can also be useful.

1. Define the circumstances under which observation will occur

This will involve users, often these will be students. Here, it is assumed that they will already be performing a predefined task in an arranged session, such as having been invited to an evaluation session for a developmental piece of software, or be participating in a class using new technology-based materials or techniques. For most evaluations, the main use of observation is as a supplemental technique to complement other standard tools. One exception when observation may be the primary tool is where it is specifically employed to provide information about a known gap in existing provision. The technicalities of setting up software evaluation sessions or courses incorporating technologies are important, but are discussed elsewhere.

2. Identify the key areas of interest

Areas you might consider include:

- ◆ How do users interact with the interface of a developmental piece of software? Are they using it in the way expected? Are they using it in innovative or unexpected ways? Are they having difficulty? In what areas, or with what tasks?
- ◆ How do students interact with the new piece of learning technology? Do they discuss the task in question? Do they develop a working strategy for the task or just plunge in? Are the instructions adequate? How long does each section take them? What is their attitude while working? How do they cope with the interface?

3. Plan your observation session

Will the observation be active or passive? Will you be able to interact with the students or users during the session or will this be intrusive?

4. Design the session proformas for the observer

These should be self-explanatory in style, and act as a prompt sheet for the observer during a session. They should include a reminder of all the key points to watch for and provide a straightforward way to quickly mark down all findings of interest plus any additional information.

It may be appropriate to design a number of different proformas for different types of session and circumstances. Or the original version could be adapted.

Provide a mechanism for identifying individual participants across different evaluation instruments, but retain the anonymity of participants unless specific permission is obtained.

Gradual changes in user behaviour may not be noticed if every session of a course is observed.

Agree the observation schedule with the person running the sessions.

Confirm and book in advance any technical help you will require, such as with video observations.

Make sure the tutor or organiser understands what you need them to tell the group. Consider giving them a checklist.

Think about how the observer should be introduced. This will affect the openness of the group towards them.

Make sure different groups observed under the same circumstances have the same proforma applied, to allow direct comparison.

Other Relevant Pages

Recipes

- Ethnography
- Focus groups
- System log data
- Resource questionnaires
- Split screen video
- Designing experiments

Information Pages

- Learning styles
- Transcribing
- Interviewing
- Student sample
- Working with groups

5. Decide when your observation sessions are going to take place

What stages of a taught course will provide the most information? For example, you may choose to cover all sessions: the first, middle, and last sessions; or to select several at random. Which are most likely to be key in students' learning or gaining technological familiarity? When will the key material be covered? Missed sessions can benefit from feedback by the tutor.

How often is it necessary to monitor? It is uncommon to need to observe every session of a course, but if there are a limited number of invited evaluation sessions, particularly if these involve different user groups, it may be beneficial to observe all.

6. Arrangements on the day

To achieve the best results:

- ◆ Make sure the tutor or session organiser explains why an observer is there. Make it clear to the participants that the reason they are being observed is to help improve the course and/or software (whatever is the focus of the evaluation), and that it is not a judgement of their skill or ability.
- ◆ In the case of courses, make sure participants know that the observation is independent, and will not affect their course grades. Emphasise anonymity for any records, published material, or subsequent reports, unless specific permission is granted by an individual participant.
- ◆ Where active observations are agreed and are practical, adopt the technique suggested in the "Split Screen Video" section. Remember that students are likely to be more open with another student or an outsider than to another member of staff from their own department.

7. Analysis and feedback

Generally a few key points immediately emerge from an observation exercise. Immediate feedback to the development team or course tutor is often most effective in these cases, as it allows direct action.

More detailed analysis should draw out the key issues for each theme, and these can be compared to the corresponding findings from using the other evaluation instruments.

Some themes will yield little or no information and can be dropped. Others can be adapted for more detailed observation. For an ongoing series of observations using the same group, such as in a course, it is often more useful to immediately adapt proformas before the next session. However, where different groups are being observed under the same circumstances, it is more valuable to keep the exact format for each group, to allow direct comparison.

Variations

- ◆ Items listed on observation proformas can be cross-referenced to specific themes addressed in other evaluation tools being used, such as questionnaires or interviews, to give different angles on the same evaluation topic.
- ◆ Proforma items can be cross-linked to original rationales for software, projects, or courses.
- ◆ Incorporate or substitute video observation, either of the participants or of the computer screen. This is particularly useful if detailed analyses of user actions are required, e.g. for HCI studies, or if a trained observer cannot attend all of the key observation sessions.
- ◆ Introduce monitoring software to record user actions during use of the software, such as mouse clicks and keyboard actions. Also useful for analysing a user's route through material or usage patterns for included software tools.

Cathy Gunn
University of Auckland,
New Zealand.

System log data is a step-by-step recording of user interaction with a software program. The level of detail recorded is determined by a purpose-built program and can include mouse clicks, menu calls, and all operations on objects. Modelled on a 'record' facility, log data is a useful way of tracking user preferences and navigational choices.

Uses – some example situations

- ◆ Any evaluation looking at how students work with courseware
- ◆ Learning evaluation: for triangulation with other data sources, e.g. questionnaire responses, interviews, focus group discussions, course grades, observation reports
- ◆ Usability evaluation: as an objective means to track use patterns, errors, navigation paths and time on task

Process involved

1. Plan

The purpose of the evaluation will determine what data needs to be logged, how it should be recorded and summarised. The evaluation plan provides the specification and level of detail for establishing the logging process. This may be 'high level', e.g. number of logins and time on task per student, or 'low level', e.g. key strokes, mouse clicks, navigation choices and frequency of errors. Logging is done by a program running on the server from which the courseware is accessed.

2. Data Collection

Raw data from the logging process is imported into a statistical package such as SPSS, StatView or Excel, then sorted into categories defined according to purpose, e.g. as individual student, group or functional profiles. Too much rather than too little data, and flexible format for output are good principles, as it is difficult to anticipate significant but unexpected results. It may be easier to deal with some redundant data than to modify the logging program during the evaluation. Statistical tests can be applied if appropriate, and emerging patterns may feedback into an iterative courseware design process and be used for triangulation purposes, or be the focus of further evaluation.

3. Reporting

In the final analysis, conclusions are drawn from or supported by system log data. Statistics are available for reporting purposes, and graphs, charts etc are easily produced. The data sets may also be used as a basis for comparison following modifications to courseware or to supporting environments. Graphs produced from log data show trends in a form that is easy to read and interpret. Statistical packages make this, and application of various tests, a simple task.

Variations of this technique

Log data can also be used to evaluate an integration strategy, i.e. do statistics show that courseware is being used as anticipated? If it is not used as expected, then other means may be employed to investigate the situation.

Log data is an objective measure which produces reliable statistical data related to a range of issues such as use patterns, usability, integration strategies and perceived usefulness of courseware. However, it does not answer 'why' questions and further evaluation is necessary to find explanations for occurring phenomena.

Careful planning is required as nothing can be captured 'after the event'. The major task is setting up the system to record the precise level of detail required.

Use statistical tests only where numbers and circumstances justify.

Other Relevant Pages

Recipes

- Questionnaires
- Interviews
- Focus groups
- Observation techniques
- Trials
- Resource questionnaires
- Checklists

Information Pages

- Statistics questions
- Pedagogical toolkit
- Student sample
- Working with groups

References

A toolkit for embedding C&IT into the curriculum

A wide variety of communications and information technologies (C&IT) are now available, and these offer education a broad range of potential benefits. However, the uptake and use of these resources is patchy at best. This mismatch between potential and use has been seen as increasingly important. However, although there are numerous incentives to increase the use of C&IT, concerns have been voiced that these may ignore issues concerning the appropriate use of the resources. This section outlines a pedagogically sound methodology for integrating C&IT into the curriculum. This builds on existing frameworks for designing courses involving C&IT and provides a context for evaluating their impact.

The pedagogic toolkit

The toolkit allows practitioners to integrate C&IT into their teaching in an appropriate manner. It is based on a number of assumptions, the most important of which is that no single teaching media is the best for all situations and styles of teaching. It is also assumed that if practitioners can adequately describe and compare teaching media, they will be able to make informed decisions about which are best suited to their particular needs.

The toolkit consists of the following key steps that guide practitioners through the process of redesigning a course:

1. Review the existing course to identify its strengths and weaknesses
2. Identify additional suitable teaching media
3. Use an elimination table to select the most appropriate teaching media for the course, considering:
 - ◆ the preparation time required (both in staff time and other resources)
 - ◆ the flexibility of delivery (in terms of the time and location of staff and students)
 - ◆ the educational interactions supported
 - ◆ the local factors that affect media use (opportunities or hindrances to the integration of C&IT)
4. Link the distinct components resulting from the steps above to form an integrated course

It should be noted that the term “teaching media” is used to cover traditional methods of teaching, such as lectures and

seminars, as well as mediated teaching such as web pages or open learning texts.

Tables are used to structure the descriptions and comparisons required by the toolkit. Each table describes one particular aspect of the media being considered, and relates to one of the elements covered by step three of the frame-work above. Practitioners using these tables are required to edit the entries for each media in order to reflect the way that they actually make use of them, rather than simply assuming that all teachers use resources in the same way.

Once comparisons have been made in terms of preparation time, flexibility and educational interactions, the resultant short-list is compared in terms of the ‘local factors’. X-opportunities or hindrances that will affect how successful the media is. This ensures that the resultant course design is sensitive to the pragmatic concerns of practitioners. A fifth and final table is then used to collect the outputs from each of the comparisons. This allows the progressive elimination of unsuitable media, and is referred to as the elimination table.

In a sense the toolkit involves decomposing and reconstructing a course. The final step in the toolkit emphasises this in the form of a ‘before’ and ‘after’ learning plan for the course. This representation has two advantages: it provides a structure for the decomposition and reintegration of the course, and it enables a quick comparison of the courses in terms of the educational interactions supported.

Conclusion

Feedback on the toolkit has highlighted its role in Quality Assurance. Inherent in its design is that the process of integrating C&IT should be structured, with justifications of decisions that have been taken. The elimination table, described above, is the clearest example of this.

The step by step guidance of the toolkit provides a supportive method of transforming and enhancing courses. Rather than being prescriptive, it seeks to highlight issues for the tutor to consider, and provides tools that can be used to address

Grainne Conole & Martin Oliver

LaTID, The Learning Centre,
University of North London

Patrick McAndrew
Heriot-Watt University.

Uses

Planning trials allows you to pace the introduction of new techniques with your students. You can decide at what stages you need to test your developments, and who can give you the right sort of feedback so that improvements can be made. Once trials are in your plan, they also become a deadline to work to; and can make all the difference between something that is really used and something that looks nice, but never gets put in front of any student.

Process

Make sure your aims are clearly defined.

1. Identify your innovation

The first stage is to recognise that you are planning a change or development that needs some feedback from others. Explicitly defining what it is that you are building and the scale of innovation involved is important. This is when you need to judge what you want to achieve with your students.

Make sure you will achieve everything you need to before your students need it.

2. Prepare a schedule

Work out a development plan. This could be your chance to try a project planner, but a simple idea of the tasks involved and content needed, matched to when you might tackle them, is all that is needed. If you make “use with your students” the only deadline in this plan, you are running a risk that it will not be ready and it will not be suitable. A series of small deadlines is easier to meet than one big one.

Include something small and early to confirm you are on the right lines.

3. Pick your time(s)

From your plan, look for an early point when you should have something available that can be used by students. You might have issues of what tools are needed and how to present information, as well as the choices about the content and teaching approach. It could be possible to address some of these questions separately and in ways that can be simply organised. However, you should also be looking for a trial of a more or less complete system before your final deadline. Note the time or times for these trials and treat them as commitments. Just identifying the trial can make a change in how people proceed with their work and the direction of effort in a project.

You might need to be flexible with your original schedule to fit in your trials and find the right users.

4. Find your users

Having decided what it is you would like to find out and when you want to do it, you need to decide who with. Trials do create work for everyone involved, so you have to value your own and your students' time. Small scale trials with friendly users can answer some initial questions, but the best groups of users will be those who are closest to the eventual users – real students. It can be worth paying a small fee to get people to spend enough time using your system, thereby giving good feedback. In early trials, numbers are less important than the quality of information.

Look for methods that will give you useful results and information you can act on.

5. Carry out the evaluation

The other sections of this book will guide you towards ways in which you can perform the evaluation itself. Picking the right evaluation method will certainly help get good quality information but there is also value in the less formal feedback from how the trial went - how difficult it was to organise, the stress levels in keeping things running, the relief when it is over.

Give yourself some time to think about the feedback and carry out any changes to your plan.

6. Use the results

Carrying out trials is not an end in itself. The feedback you have from your users needs to be used in your development process. Make some initial judgements about how well it went – consider if there needs to be significant changes made to your plans; there might be some radical suggestions that can be identified. More likely you will have some generally positive information together with some complaints or weaknesses that have been identified. Try not to let these dishearten you, but instead see that you now have a clear basis for action to address these. You might need to weigh up some insightful comments against consistent patterns across your users. If the trial was local you can always ask people to clarify their points.

7. Do it all again!

Trials are part of the process of formative evaluation; this means that you are making changes to what you are offering and, of course, to validate those changes you really should carry out some trials and evaluation.

Variations

Trials need not just be used with new material. The process of integration of any innovation with the rest of the course is a very important process. This means that a trial can extend to incorporate existing activity, or that a trial can be planned to find out information from students before they are introduced to anything new. This can identify their wishes and current problems so that you can be more secure about addressing their real needs.

Other Relevant Pages

Recipes

- System log data
- Observation techniques
- Checklists
- Focus groups
- Interviews
- Cost effectiveness
- Nominal group techniques
- Pre and post testing

Information Pages

- Isolation or integration
- Pedagogical toolkit
- Pre and post testing
- Questionnaires
- Student sample
- Statistics questions
- Working with groups

Evaluation report outline

Serving suggestions

References

Avoid cancelling trials but do value the time and involvement of learners. If you are not at the stage you hoped you would be, think about refocusing the trials and scaling them back rather than not carrying them out. There is always valuable data to be had.

Keep track of developments. Trials give a co-ordination or management role for the evaluator as they can provide the main means to keep track of the development work and ensure that systematic testing is continually given emphasis. This is rather different to the "traditional" view where the evaluator is there simply to gather data at the end that will prove it works. Other sections in this publication will suggest methods that can be applied but using a formal evaluation framework will often help all forms of feedback from a trial.

A quick and easy recipe for chocolate cake. Minimal culinary skills are required and no baking.

Uses

- ◆ Celebrations: birthdays, successful funding bids, end of term
- ◆ Bribery: colleagues, students
- ◆ Reward: self

Process

1. Planning

Assemble the following ingredients and equipment:

Metric to Imperial Conversion

25 g = 1 oz
2.5 cm = 1 inch

250 g digestive biscuits	food processor
125 g soft brown sugar	saucepan
125 g butter	wooden spoon
50 g raisins	18 cm sandwich tin (greased)
3 tablespoons cocoa powder	fridge
1 egg, beaten	knife
few drops vanilla essence	

2. Identifying your population sample

Consider carefully the population you are making the cake for:

- ◆ Do they like chocolate?
- ◆ Do they have any allergies or special dietary requirements e.g. are they vegans?
- ◆ Will they be sufficiently appreciative?

Make sure you select an appropriate population sample size to ensure satisfactory results. Too small a population and your cake will be wasted; too large and there will not be enough to go around, resulting in bad temper and your own unpopularity.

3. Design

- ◆ Roughly chop the biscuits.
- ◆ Place the sugar and butter in the saucepan and gently heat until melted.
- ◆ Add the raisins and cocoa and then remove from the heat.
- ◆ Add the egg and vanilla essence.
- ◆ Beat the mixture (by hand or in the food processor) until well mixed.
- ◆ Turn into the greased sandwich tin and spread evenly throughout the tin.
- ◆ Chill in the refrigerator until set.

Attractive and well-timed presentation will enhance the uptake of the cake

4. Delivery

Cut into wedges, arrange on a plate and serve.

5. Analysis

Energy per 100g

KCals: Excessive

Variation

The cake can be made with any plain biscuits.

Nuts, glacier cherries may be added to the mixture.

A drop of whisky gives the cake a uniquely Scottish flavour.

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.

This summary is adopted from the work of Morris, L.L. (1987) "How to Communicate Evaluation Findings." and is a useful guideline for anyone having to write an evaluation report. Too often we rush into the actual evaluation without giving due consideration into how we are going to communicate our findings once the evaluation is complete. This framework is also useful when planning an evaluation as it covers all the areas that could potentially be involved in having to conduct one.

Section 1 - Summary

Make this a short summary for people who won't read the whole report. Give the reasons why the evaluation was conducted and who it is targeted at together with any conclusions and recommendations.

Should cover:

- ◆ What was evaluated?
- ◆ Why the evaluation was conducted?
- ◆ What are the major findings and recommendations?
- ◆ Who is the report aimed at?
- ◆ Where there any major restrictions placed on the evaluation? and by whom?

Section 2 - Background

In this part, cover the background to the evaluation and what is was meant to achieve. The program should be described and the depth of description will depend on whether the intended audience have any knowledge of the program or not. Don't assume that everybody will know. Don't leave things out but at the same time don't burden them with detail.

Should cover:

- ◆ origin of the program,
- ◆ aims of the program,
- ◆ participants in the program,
- ◆ characteristics of the materials,
- ◆ staff involved in the program.

Section 3 - Description of the Evaluation

This covers why the evaluation was conducted and what it was and was not intended to accomplish. State the methodology and any relevant technical information such as how the data was collected and what evaluation tools were used.

Should cover:

- ◆ purposes of the evaluation,
- ◆ evaluation design,
- ◆ outcome measures
 - instruments used,
 - data collection procedures,
- ◆ implementation measures.

Section 4 - Results

This will cover the results of the work from section 3 and can be supplemented by any other evidence collected. Try to use graphics (charts, tables etc.) to illustrate the information but use them sparingly to increase their effectiveness.

Should cover:

- ◆ results of the study
 - how many participants took any tests?
 - what were the results of the tests?
 - if there was a comparative group, how do they compare?
 - are any differences statistically significant.
 - if no control group, did performance change from test to test?

Section 5 - Discussion

This should discuss your findings and your interpretation of them. Always interpret your results in terms of the your stated goals.

This section should cover the interpretation of all the results in section 4. If the evaluation is not a large one then sections 4 and 5 could be combined. The results should always be related back to the purpose of the evaluation, something that doesn't always happen. Don't forget the unexpected results as they can often be the most interesting.

Should cover:

- ◆ are there alternative explanations to the results from the data?
- ◆ are these results generalisable?
- ◆ what were the strengths and weaknesses of the intervention?
- ◆ are certain parts of the program better received by certain groups?
- ◆ are any results related to certain attitudes or learner characteristics?
- ◆ were there any unexpected results?

Section 6 - Costs and Benefits

This is an optional section and would only be included if this had been part of the evaluation plan. As there is no definitive approach to investigating this whole area there will be a need to justify the approach taken. Not many evaluations look at costs but there is a growing need to include some information about this area. Evaluations and program interventions don't happen for free.

Should cover:

- ◆ what was the method used to calculate costs and effects/benefits?
 - how were costs and outcomes defined?
- ◆ what costs were associated with the program?
 - how were costs distributed (e.g. start-up costs, operating costs etc.)?
 - where there any hidden costs (e.g. in-kind contributions)?
- ◆ what benefits were associated with the program?
- ◆ what were measures of effectiveness (test scores; program completion etc.)?
- ◆ were there any unexpected benefits?

Section 7 - Conclusions

This section can be the most important section in the report apart from the summary. Some people will only read the summary and the conclusion section. Conclusions and recommendations should be stated clearly and precisely and these might be presented as a list as readers can easily scan them. Don't expect everyone to read your report from cover to cover. Make sure that you get your main points across in the opening summary and in the conclusion.

Should cover:

- ◆ what are the major conclusions of the evaluation?
- ◆ how sure are you of the conclusions?
- ◆ are all the results reliable?
- ◆ what are the recommendations regarding the program?
- ◆ can any predictions or hypotheses be put forward?
- ◆ are there any recommendations as to future evaluations?

Philip Crompton

Research Fellow, Institute for Education,
University of Stirling.

- 1 Checklists
- 2 Concept Maps
- 3 Confidence logs and questionnaires
- 4 Ethnography including observational techniques
- 5 Focus groups, user logs and post-trial questionnaires
- 6 Nominal Group technique
- 7 Pre and Post testing and tracking log data
- 8 Questionnaires and Interviews
- 9 Planning trials

Checklists

Colin Milligan
TALiSMAN, Heriot-Watt University

Background

TALiSMAN has used electronic feedback forms to record participant views on our Online Course 'Using the WWW in Teaching and Learning', which has been delivered four times to approximately 300 staff at Scottish HEIs.

Methodology

Forms were presented for completion over the WWW and responses collected by email, enabling entirely paperless administration. The forms were a mixture of open and closed questions. These online evaluation forms were used to check on the usability of discussion forums, to identify which parts of the course were of most use (and of least use) and ultimately to revise the course itself. Unfortunately, as completion of the forms was voluntary, this meant that the response rate was low.

By asking the participants whether they felt they had met course objectives, we were able to measure whether the course materials were matched to the participants - a check for prior knowledge was also made to check that the course was being properly targeted.

Reflections on the study

These evaluations provided a quick and efficient way of collecting information about the participants perception of our course. The checklist format was especially useful in providing a large amount of low level information about many aspects of course content, design and delivery.

Concept Maps

Judith George
Open University

Where and when was the study carried out?

I teach an Open University second level Arts course on the Homeric epics. Students explore the two texts from several angles - oral poetry, war poetry, and the archaeological context of the Homeric period. They come into the course from a range of backgrounds. Some have been fired by archaeological programmes on the TV or visits to ancient sites in Greece to find out more about early Greek culture; others may be coming in with little or no experience or understanding of archaeology, because they are fascinated by the literature. Learning is primarily through the course material, and the assignments, with only a scatter of tutorials throughout the year. So tutorial contact time is therefore valuable as a scarce commodity, and must be used carefully to match and meet student need as closely as possible.

I have usually run a tutorial at the beginning of their work on the archaeological material, pitching it at the explicit challenge in the course text - to explore the differences between the disciplines of history, literature and archaeology, and what each can validly contribute to the others and to our understanding of the poems and the world in which they were created.

I had introduced the habit of ending each tutorial with a concept map exercise from the start, and ended this tutorial similarly. Because of the scarcity of tutorials, I had had to make the first occasion an experiment, which they would have to find immediately useful and convincing if we were not to drop the idea for future sessions.

How many staff and students were involved?

This evaluation just involved myself, as tutor, and the 13 students who attended the class, and took about 15 mins at the end of the tutorial.

What were the aims and objectives of the evaluation?

To gain insight into cognitive aspects of the tutorial session; had the students understood the conceptually difficult differences in academic discipline and methodology here? Where there areas which they found more difficult than others, and would need further work? Did the approach I had adopted sit well with the way in which they were working with the texts and the course materials?

Both the students and I found that this put the design of the tutorial on our joint agenda. They saw that a tutorial didn't just happen - it was carefully thought out to meet their needs, and it could do that better if they thought more about what those needs were and told me of them. They also gained confidence in my genuine wish to help them, and in the better progress which they could make if we talked together about the process of learning, as well as its content.

What did I find out?

That a significant number of the students had not even got near engaging with the conceptual focus of the session, because they were ignorant of the basics of archaeological methodology. They had not really understood the course material on the subject and were still struggling with concepts such as stratification, dating techniques and so on. This had not been apparent from discussion in class, because there were a few archaeological enthusiasts who had led either led the discussion or contributed helpfully and so on. The rest, from the feedback, had grasped what they were saying superficially, but, given their unease with the more basic material, could not engage with the ideas on their own and could not apply what knowledge they had.

Confidence Logs and Questionnaires

Helyn Thornbury
University of Strathclyde

Where and when was the study carried out?

The evaluation study involved a class studying Manufacturing Systems during the year 1997/98. This was a one semester course which aimed to cover a basic introduction to manufacturing systems, including stock management techniques. Previously to this academic year, the class had been taught using a traditional combination of lectures, tutorials and labs. The course lecturer produced 2 computer based learning modules which cover the equivalent material to 4 hours of lecturing. These modules were integrated into the class and the related lectures removed from the class programme. An introductory lab session was added to introduce the system to the students. The students were also given a notes template to accompany the modules.

What are my reflections on this study?

The concept maps are a technique which is easily operated even within the strict time limits of an Open University tutorial on a 30 point course. I gained a valuable insight into the extent of students' understanding and the coverage of learning in the tutorial, which I could not have gained from observation of behaviour; and also, importantly, into the reason for shortcomings in understanding and progress. On the basis of that insight, I could simply do a little remedial work with this year's students to bridge this gap, and redesign the tutorial for the next year's group to start with some diagnostic work to identify the extent of those particular students' understanding of archaeology, thus being able to start work at a point appropriate for them.

The students explicitly welcomed this addition to the tutorial programme. In practice, they found it an invaluable period of "time out" to reflect on what they had covered, to consolidate, and to identify gaps which could be filled quickly, or more major problems which we could then plan together to address.

How many students were involved?

The class is composed of several distinct groups of students: Business, Engineering 1 and Engineering 2. Ninety-three students were involved in total.

What were the aims and objectives of the study?

The evaluation aimed to investigate several areas. These were:

- ◆ Usage of the CBL materials
- ◆ Level/depth of content covered
- ◆ Student attitudes to the use CBL
- ◆ Appropriateness of organisation of the class (e.g. group size)
- ◆ Notes template
- ◆ Learning development of the class.

Which evaluation techniques were used?

Two data collection techniques were used in this study: questionnaires and confidence logs.

Two questionnaires were designed for the study: The first was very short, concentrating on students' previous experience with computers and CBL materials, attitudes to CBL and their initial reactions to the system. The students were given this after their first use of the system during the supervised labs. This was intended to correspond with the steepest part of the learning curve, before the students had become familiarised with the interface and system design.

The second questionnaire was given to the students after they had finished using the system. This was more substantial than the first and covered a range of areas, not only those relating to the modules but those involving wider organisational and support issues for the class.

In addition to the questionnaires, the students were asked to fill in 2 confidence logs - the first before they had used the material, to provide a baseline of information. The second after several weeks use to see any development in their confidence during this period. In this case, the confidence logs consisted of 8 statements based on the knowledge and skills objectives of the course. Information from the confidence logs gives an indication of the change in confidence in these areas during the time investigated. Of course, any changes cannot be tied specifically to the use of the modules but it does show the effect of this part of the class as a whole.

What did we find out?

The response rate for the investigation as a whole was very high. The following percentage return rates were achieved:

	%
Questionnaires:	
First	74
Second	70
Confidence logs:	
Baseline	93
Second	81

The results of this evaluation indicated that the new resources combined effectively with the existing material, as expected. An interesting result was the difference between the different categories of students. The confidence logs revealed previous knowledge of particular parts of the course by the business students of which the lecturer was unaware. Further investigation revealed an overlap in material with a second year class for the business students.

What were my reflections on the study?

The evaluation highlighted parts of the material which the three different groups in the class found challenging, which were not the same areas! This information allowed the lecturer to further tailor the class to the specific mix of students.

Ethnography including observational techniques

Chris Jones
Liverpool John Moores University

Where and when was the study carried out?

This research was undertaken at Manchester Metropolitan University between 1994 and 1996. The course unit observed was part of the BA in Information Technology and Society Degree. The course "Technology in Communications" was a second year option taught on-line using the FirstClass computer conferencing system. The approach was to deliver the course 'online' as much as possible. The course unit is still being run in a slightly modified form, but it is now likely to migrate to the First Class intranet server which will provide an internet gateway.

What were the aims and objectives of the study?

The evaluation was 'illuminative' in that the question asked was 'what happened when a course with an explicitly collaborative aim was taught using a conferencing system as the means of delivery?'. An ethnographic methodology was employed to generate an adequate description of 'just what' happened when university education was transposed from a traditional setting into the new technology.

How many students were involved in the study?

Manchester Metropolitan University (MMU) is situated close to Manchester City centre. The University recruits students from a wide range of academic and social backgrounds and has a large number of 76 home students from Greater Manchester and the surrounding region. The course ran in 1994/5 with 8 registered students. The students were mainly native English speakers five of whom were male, and three were female. In the second year, the registration included a large number of students whose first language was not English. The total registration for the course was 20 students, 7 female and 13 male. The English language competence and experience of this group varied considerably, especially their written fluency.

What were the main findings?

During the two years of observation, the work that was done was organised in part on-line, though much of the activity was off-line and ratified rather than conducted on the system. Students when working amongst themselves, distributing their work between on and off-line working. Work carried out in the system, without physical contact, was often a case of division of labour. The system would be used to allocate tasks to individuals writing sections and later a named individual would take the sections and combine them into a final text. Often there was little or no consideration given to re-drafting. These submissions were essentially composite, more or less coherent depending upon the skill of the student compiling the final copy.

The expectation for the course unit was that it would involve the development of collaborative skills through group working and the development of peer interaction. Students were observed orienting themselves to course requirements so that the conference transcripts reflected the students and tutors understanding of what they ought to do. It was found that the transcript served as an official document. In many instances, on-line activity was an artificial construct consciously produced as material for assessment.

Interestingly, collaboration was not achieved in the expected form: students did collaborate, they were observed off-line working in groups talking over and around machines. The work would then be presented in the conference as individual messages, a division of labour. However, the expectation, of peer interaction was that students would draft and redraft work on the system and that the process of collaboration would be visible, did not happen.

Reflections on the study

The research cast doubt on the degree of collaboration and the character of the collaboration that took place. It also questioned the reliability of the use of transcripts for research and assessment purposes. Transcripts were often incomplete and only provided partial records of the activity of the conference. On-line text was often a public display. Therefore, the transcribed record was not so much about what happened but more a reflection of what participants thought ought to be happening. Students and tutors recorded those things they believed were required to fulfil the requirements of the course.

Ethnographic evaluation, by adopting an 'indifferent' stance to the course and observing informally the practice of students' and tutor's on and off-line, provided information not readily accessible by other methods of evaluation. It revealed the detailed way in which students co-operated together in order to achieve what they took to be the course requirements. A more formal evaluation procedure may have concentrated on the 'learning process' and not have considered some of the observed interaction might relevant to education.

Focus Groups, User logs and Questionnaires

Erica McAteer and Liz Leonard
Glasgow University

Where and when was the study carried out?

The initial study was carried out during 1996/97 within the Institute of Biomedical Life Sciences at the University of Glasgow and was part of the third year of the University's development of on-line formative and summative assessment materials, produced using Question Mark. A considerable amount of work has been carried out since then and both development and evaluation processes are still ongoing.

How many staff and students were involved?

The main work has been with Level One Biology, the first year course which all students take before selecting from a menu of options in Level Two towards specialism in third and fourth year. The class takes in around 700 students each year and provides four subject modules and a study project module. Four Associate Lecturers carry the main load of laboratory demonstration, teaching management and administration, with 19 module lecturers providing the learning content through readings, practical work, lectures and assignments.

Which evaluation techniques were involved?

During the first year of development, a set of self assessment items was produced for just one of the subject modules, 'Molecules, Cells and Genes'. The items were authored by the subject lecturer in collaboration with an associate lecturer with experience in using the software and with expertise in the field of assessment, and in objective testing in particular. The module ran in term one, the tests were trialled by student volunteers (220) after a revision lecture in term two, one or two weeks before the module exam.

User logs, post-trial questionnaires and an item in the standard post course evaluation questionnaire provided some information (in the main, positive) about the pilot. A focus meeting was held a few days after the exam, with six students at the end of an afternoon lab. class and refreshments were provided.

Following this, further tests were developed for modules held in the second half of the year, and these again evaluated by students through logs, questionnaires and two focus meetings. For 1997/98 students were provided with self-assessment tests for all four subject modules. Further focus group meetings, concentrating on different issues to do with item design, interactivity, timing and addressing issues of on-line summative assessment were held. Development continues through this (1998/99) session and tests for second year modules are under design and pilot.

What were the aims and objectives of the evaluation?

- ◆ To establish student use of and attitudes to on-line assessment materials.
- ◆ To test the authoring software against our need to improve practice in objective test provision.
- ◆ To inform lecturing staff about the value of this form of assessment, in terms of development effort and pedagogical and practical benefit.
- ◆ To advise institutional management about investment issues - software, support staff, equipment resources.
- ◆ To provide the grounds for reflective development of assessment practice.

What did we find out?

Things that we expected, things that we didn't expect.

Students welcomed the tests as self-assessment and revision resources.

They particularly valued immediate and, where suitable, directive feedback.

The reasons they gave for their judgements reflected concerns beyond the practical in that they felt that the tests not only 'helped them know where they were' but also 'gave a better understanding of the course content'. It was the strength of their feeling that all modules should have such resources that moved development forward earlier than planned.

They picked up differences in question style and rhetoric, confirming our expectation (hope?) that the interactivity enabled by the software, and the potential for 'deeper' learning to be addressed, would be perceived by them. It was also welcomed by them.

The content of their discussion also indicated that attitudes to such uses of computer resources were shifting towards acceptance as familiar and commonplace elements of the classroom. That more than half the students said that they would have no objections to being summatively assessed in this way was a surprise. Because of the richer feedback provided by the method, allowing argument and elaboration as part of the data, we realised that what objections there were often had more to do with objective testing itself, rather than computer based assessment. This echoed staff feeling closely, and was important for the design and development of the overall assessment procedures for the modules and the course as a whole.

What are our reflections on this study?

One of the most interesting outcomes of the study, from our point of view at least, was the change in the staff attitude. Having seen examples of the kind of objective testing that could be supported and the serious and reflective nature of the student feedback, the lecturers began to realise the potential of the system. The further development of the new assessment procedures depended upon such realisation.

Rather than relying solely on the quantitative feedback from logs and questionnaires, or the more qualitative feedback from the few open question responses received from the questionnaire administration, we were able to 'play back' the transcripts of the focus meetings to the staff concerned. We felt that they would be the best interpreters of such feedback.

The methodology itself has now become an integral part of the long-term development of assessment procedures within the Level One class, and is becoming so for Level Two.

Nominal group technique

Gaye Manwaring
Northern College

Two quick tasters:

1. We used a nominal group technique to get views about the use of FirstClass computer conferencing for tutorial support and peer interaction by a group of mature distance learners. The evaluation was carried out while they were together for a face-to-face workshop and revealed a variety of views. Some were about the content of the material, others about the learning process and others about hardware problems.

2. We used a quality/importance grid as part of the review of an innovative course. Some issues of high importance and low quality were beyond the immediate control of the lecturers but were having a damaging impact on the course. Some students were denied access to computers at their workplace and could not complete assignments on time. Once alerted to the problem, the course director was able to approach the employers and negotiate better support for course members in the workplace.

Pre and Post testing and tracking log data

Robert Newton
The Robert Gordon University

Where and when was the study carried out?

The model described was used with an undergraduate group of Librarianship and Information Studies students at the Robert Gordon University. The students were in their second year of a four year degree programme and initial testing involved creation of the student profiles as described in the recipe and in addition, input on general academic performance and attitudes from the previous year of their studies. The Gregorc learning style delineator was used to determine learning style.

How many staff and students were involved?

Only 20 students were involved in the test. The students were given an introductory session on using the CAL package and no student encountered particular problems with this. Although an examination of the tracking mechanism built into the programme showed some interesting variations on how students used the package.

What were the findings?

Analysis of the results from pre and post testing in both directions (CAL package then lecture and lecture then CAL package) showed that there was certainly a variation in performance level between students as expected, both showing some improvement. However, the level of improvement varied considerably between individual students. There appeared to be some evidence that this was linked to learning style and this was partly confirmed by an analysis of the tracking logs from CAL sessions that showed marked differences in approach to using the CAL materials. Variations in performance related to personal factors, motivational factors and previous academic performance (including facility with computers) were not seen as particularly significant. Numbers involved were too small to show statistically significant variations. The larger tests currently being conducted using a postgraduate cohort of 65 students are expected to yield more useful data.

Questionnaires and Group Interviews

Judy Goldfinch and Kirsty Davidson
Napier University

Where and when was the study carried out?

During 1997 and 1998, a suite of computer-based learning software (Mathwise and Statwise) was used by several Scottish universities' mathematics departments. This was part of the SHEFC funded SUMSMAN project which saw all the Scottish university mathematics departments sharing resources and materials for (mainly) first year service teaching.

Departments at different universities integrated the software into their teaching in different ways, and, indeed, integration methods varied between members of staff at the same university. Some got the students to work through the package in supervised lab sessions either after or, in some cases, before meeting the topic in lectures. Some used the package in a similar way but did not staff the labs. Some replaced conventional lectures on a topic with directed study of part of the package. Some used the software as a remedial aid only, or purely as an optional extra resource whose existence was pointed out to students.

The software itself had already been extensively evaluated, but this study was designed to evaluate the various ways of integrating it into courses. Both staff and student reactions were sought, but only the aspects involving students are reported here.

What methods did we use?

We decided to survey all students involved in the study, by means of questionnaires. Although these included open questions to allow for unanticipated reactions, most of the questions were closed. This meant that we could gather and analyse large quantities of data quickly and cheaply and would be able to spot patterns and gauge the extent of variety in student views. The questionnaire design and analysis package Sphinx was used to handle and analyse such a large amount of data.

To gain a deeper insight into student views, two structured group interviews were held with students. To encourage comparison of integration methods and to spark discussion, each of these was planned to involve a group of about ten students from each of two universities who had studied the same topic or topics. The sessions were run by video-conference so that no one had to travel.

How many staff and students were involved?

In all, several hundred students and nearly twenty staff from five universities were exposed to the software in some way. Only three universities and about fifteen staff actually issued questionnaires to their students but nearly three hundred were returned, covering all the integration methods used and students from a wide range of faculties (arts, science, business and engineering!). Staff merely distributed and collected the questionnaires, all the analysis was done by one individual. Luckily, using the Sphinx software package made this relatively painless.

We expected ten or so students from each university to attend the two discussion interviews. However, on one occasion only five turned up at each site and on the other occasion no students could be persuaded to attend at all from one site. A member of the staff involved in the study was present at each site to lead and record the discussion, as well as to run the short ice-breaking session held prior to the actual interview.

What were the aims and objectives of the evaluation?

- ◆ To determine student reactions to the various ways of using the package, and to identify the 'best' ways of using it.
- ◆ To identify factors influencing how useful students found the package.

What did we find out?

The responses were remarkably consistent. Students overwhelmingly thought that computer-based learning was:

- ◆ best used in staffed labs
- ◆ most useful for:
 - practising material that they had just been taught,
 - trying to understand material that was unclear in a lecture,
 - revising for an assessment.

Factors found to influence how useful students found the package were:

- ◆ the imminence of an assessment,
- ◆ the enthusiasm of the staff and how well they sold it to their students,
- ◆ how well the content was perceived as matching their lectures,
- ◆ the speed at which the equipment allowed it to run,
- ◆ the amount of feedback provided by the package, and
- ◆ provision of accompanying paper-based support material.

Surprisingly, students who found the pace of the lectures 'too fast for me' were less likely than other students to find the CBL software useful.

What are our reflections on the evaluation methods used in this study?

We feel that the evaluation methods were very successful and achieved our aims.

The questionnaires yielded information in bulk and provided reassurance as to the reliability of the information. Despite being quite long (4 pages), there was little evidence of lack of care in their completion. Almost all students were still willing to write sensible and sometimes extensive comments in the open questions ('Main drawbacks', 'Main benefits') at the end. This was helped, as always in questionnaires, by making the early part of the questionnaire quick and easy to complete, and maintaining an interesting and logical progression of questions.

By getting students to write their matriculation number on the questionnaire, further analysis of how responses compare to, say, student ability or success in the course can be carried out later.

The group interviews gave fuller information and being a group session allowed comments to be taken up by other people and explored from several angles. Five students and one staff member at each site actually worked very well with all students feeling that they should say something. Having students from two different groups was also very successful with fresh ideas and differing experiences providing interest and allowing discussion of new points that had not occurred to one or other group. Using video-conferencing was a bit intimidating at first but people soon forgot about it. Staff involved felt that the benefits of the two-group sessions out-weighed any hassle in setting up the video-conferences.

Having the software at hand to refer to was helpful, and a demonstration at the start helped jog memories and break the ice. An initial ice-breaking exercise helped as well, as would have, in retrospect, provision of name badges. Forty-five minutes was found to be ample time for a session.

In this study the interviews were conducted after at least some of the questionnaire data was studied which allowed queries raised by the questionnaires to be investigated more fully in the interviews.

Planning trials involving use of questionnaires and focus groups

Patrick McAndrew
Heriot-Watt university

Where and when was this study carried out?

This was part of a fairly large European project funded to develop both new network based technology and material that used innovative approaches to teaching. Trials were devised that involved students at one UK University to carry out tasks aimed at learning a second language (French). These students were directly recruited to take part in the trials although many were also undertaking optional modules in French. The trials took place over a period of six weeks and were scheduled in advance.

How many staff and students involved?

The trials involved twenty-four students, one evaluator, three members of teaching staff, and various support staff including some from European partners (this is not typical!).

Which evaluation techniques were used?

In addition to planning the trials, the evaluation used questionnaires that gathered the students' attitude to the technology, and confidence logs before and after the trial to provide for self-evaluation of progress. Many sessions were video taped to provide later study and some of the students were interviewed individually and in small focus groups. Students also carried out the same tasks without using computer support. Statistical analysis was used to present the data and search for significant results.

What were the aims and objectives of the study?

The main aim of the trial was to demonstrate a working and viable system. Further objectives were to see if we could measure any improvement in language ability and to determine the students attitude to the different components used in the system. In particular, high-quality desk top video conferencing was a major component of the system and we were interested in the value that the students placed on this.

What were the findings?

The trials supplied information about the relative value of aspects of the system but could not provide an absolute measure of success through comparisons. This was as expected and the value of such trials is perhaps first as project management, second as direct feedback, and last as good data for academic papers!

What were your reflections on this study?

The devising of tasks for the trials had a direct effect on the materials and technology development as anything that was to be trialed definitely had to be supported by the system. The main trials were across distributed sites and the installation process went right up to the start of the trials. This demonstrates the difficulty of completing anything before a real deadline and the power of a trial to provide that deadline. Contingency plans were in place to investigate only some elements of the complete system but in the end these were not necessary. Without these trials it is possible that this integration would never have happened.

General references relating to evaluation:

Angelo, T.A. and Cross, K.P. (1993) Classroom Assessment Techniques. Second Edition. San Francisco: Jossey-Bass.

Breakwell, G. and Millward, L. (1995) Basic Evaluation Methods. Leicester: BPS Books.

Dix, A., Findlay, J., Abowd, G. and Beale, R. (1993) Human-Computer Interaction, London. Prentice Hall.

Draper, S.W., Brown, M.I., Edgerton, E., Henderson, F.P., McAteer, E., Smith, E.D. and Watt, H.D., (1994) 'Observing and Measuring the performance of Educational Technology. The TILT project', University of Glasgow.

Draper, S.W., Brown, M.I., Henderson, F.P., McAteer, E. (1996) Computers in Education 26, 1-3, p17-32.

Further Education Unit. (1991) Towards an Educational Audit. London: FEU

Gibbs, G. (1995) Improving Student Learning through Assessment and Evaluation. Ed. Gibbs, G. The Oxford Centre for Staff Development. ISBN 1 873576 43 9

Gunn, C. (1996) CAL evaluation; What questions are being answered? Computers and Education 27, 1-3, p17-32

Harris, D., and Bell, C. (1994) Evaluating and Assessing for Learning. Second Edition. London: Kogan Page.

Hounsell, D., Tait, H. and Day, K. (1997) Feedback on courses and programmes of study: A Handbook. UCoSDA. ISBN 0 9523956 7 3

Mason, R. (1995) Evaluating Technology-Based Learning. In Collis, B. & Davies, G. (Eds) Innovative Adult Learning with Innovative Technologies, p191-199. Elsevier Science B.V. Holland.

Milne, J., and Heath, S. (1997) Evaluation Handbook for successful CAL courseware development. MERTaL publications. ISBN 1 873154 48 8

Morrison, K. (1993) Planning and Accomplishing School Centred Evaluation. Dereham: Peter Francis Publishers.

Nicholson, A.H.S. (1997) CERT: A courseware evaluation and Review Tool Published by CTI Accounting, Finance and Management.

Pratt, D. (1994) Curriculum Planning. Orlando: Harcourt Brace.

Preece, J., Rogers, Y., Sharp, H., Benyon, D., Holland, S. and Carey, T. (1994) Human-Computer Interaction. Addison-Wesley. UK.

Race, P. and Brown, S., (1998) Learning from student feedback in 'The Lecturer's Toolkit' Kogan Page. ISBN 0 7494 2496 6

Reeves, T.C., (1991) Ten Commandments for the evaluation of Interactive multi-media in Higher Education, Journal of Computing in Higher Education 2, 2, p84-113

Smith, C. and Mayes, T. (1993) 'Telematics Applications for Education and Training USABILITY GUIDE, Version 2, DGXIII 3/c', Commission of the European Communities, Brussels.

Soler, J. and Lonsberg, M. (1996) Rethinking the teaching of a University Course. Active Learning 5, p3-5, December.

Tessmer, M. (1993) Planning and Conducting Formative Evaluations. Kogan Page. London.

Vandervelden, G. (1996) Evaluation of courseware – A Practical Approach. Teaching and Learning Technology Support Network University of Kent publication.

Useful educational references including sections relating to evaluation

Ellington, H., Percival, F. and Race, P. (1993) The Handbook of Educational Technology. Kogan Page, London.

Laurillard, D. (1993) Rethinking University Teaching – a framework for the effective use of educational technology.

Ramsden, P. (1992) Learning to Teach in Higher Education. Routledge.

Rowntree, D. (1982) Educational Technology in Curriculum Development. Atheneum Press Ltd. Newcastle upon Tyne.

Rominzowski, A.J. (1988) The selection and use of Instructional media. Kogan Page.

Evaluation on the World Wide Web

ERIC Clearinghouse on Assessment and Evaluation:
<http://ericae.net/>

The ERIC Clearinghouse on Assessment and Evaluation seeks to provide 1) balanced information concerning educational assessment and 2) resources to encourage responsible test use. Includes a directory of "Assessment & Evaluation on the Net".

ELT: Evaluating Learning Technologies:
<http://www.unl.ac.uk/latid/elt/>

ELT is based at The University of North London's LaTID (Learning and Teaching Innovation and Development), and seeks to develop methodologies for evaluating, and carry out evaluations of learning technology with the aim to producing an evaluation toolkit for use by staff without experience of evaluation to use on their own courses. ELT has also developed a pedagogical framework for embedding technology into courses.

EVA evaluation tools:
http://bengt2.citu.lu.se/eval/index_eng.html

EVA is a tool for quality assurance and useful in the process of investigating the outcome of processes, courses, educational programs and creating forms for many purposes. 'EVA can also be useful in the collation of the answers!'

EvNET:
<http://socserv2.mcmaster.ca/srnet/evnet.htm>

EvNet is a Canadian network comprising membership from 14 universities, 5 community colleges, 3 schools, 21 private corporations, 5 government agencies, and 10 non-profit organisations and is involved in evaluating the effectiveness of computer-mediated communications in the delivery of education and training.

Extracts from Publications of the Oxford Centre for Staff and Learning Development (OCSLD):
<http://www.lgu.ac.uk/deliberations/ocsd-pubs/>

Extracts from: Improving Student Learning - Theory and Practice (1994) Ed.: Gibbs Improving Student Learning - Through Assessment and Evaluation (1995) Ed. Gibbs Course Design for Resource Based Learning (various subjects, various editors, 1994)

INUSE project:
<http://info.lut.ac.uk/research/husat/inuse/>

INUSE (Information Engineering Usability Support Centres) are a network throughout Europe to assist both companies and projects within the EU Telematics Applications Programme by providing usability tools and methods. The support centres promote best practice on usability, quality assurance guidelines for user-centred design, and HCI standards.

NCODE evaluation links:
<http://cedir.uow.edu.au/NCODE/info/evaluation.html>

A collection of WWW links relating to evaluation and quality produced for the National Council on Open and Distance Learning Resource based learning site.

Open University pages:
<http://www-iet.open.ac.uk/iet/PLUM/contents.html>

Open University page about methods and practice of formative and summative evaluation of multimedia materials.

RESPECT:
<http://info.lut.ac.uk/research/husat/respect/>

A collection of methods and approaches for user-based requirements analysis from the RESPECT European Usability Support Centre (part of the INUSE project). The broad aim of these methods is to elicit from users of software what they want, and feeding this information into the software development cycle; however many of the methods are also useful in the context of selecting software or evaluating implementations.

Serco Usability Services:
<http://www.usability.serco.com/index.html>

Provide a commercial evaluation service for software and web design including consultancy, evaluation tools, training and laboratory services. Many relevant publications are available from their web site.

TILT Evaluation group pages:
<http://www.elec.gla.ac.uk/TILT/E-Eval.html>

References about the development and application of evaluation methods from the institutional TLTP project at Glasgow.

Checklists

Draper, S. and Oatley, K. (1991) Highly interactive visual interfaces. Programme 2: Assessing Interaction. University of Glasgow.

Oppenheim, A.N. (1966) Questionnaire design and attitude measurement. Heinemann.

Tergan, S.O. (1998) Checklists for the Evaluation of Educational Software: Critical Review and Prospect. 35,1, p9 -18. Innovations in Education and Training International.

Concept Maps

Buzan, T. (1995) Use your head. Updated BBC version. ISBN: 056337103X

Inspiration software produced by iANSYST training products enables you to computer generate concept maps/ flow charts/ diagrams etc. Available from iANSYST, The White House, 72 Fen Road, Cambridge, CB4 1UN Tel. 01223 42 01 01 or a demonstration version can be downloaded from <http://www.dyslexic.com>

Confidence logs

Draper, S.W., Brown, M.I., Edgerton, E., Henderson, F.P., McAteer, E., Smith, E.D. and Watt, H.D. (1994) Observing and Measuring the performance of Educational Technology TILT project, University of Glasgow.

Heller, R. (1991) Evaluating software: a review of the options. Computers in Education, Vol. 17, No. 4, p285-291.

Cost-Effectiveness

Crompton, P., (1997) Calculating Cost-effectiveness of IT in Higher Education: Do the numbers all add up? Paper presented at ALT-C 1997.

Doughty, G., Arnold, S., Barr, N., Brown, M., Creanor, L., Donnelly, P., Draper, S., Duffy, C., Durndell, H., Harrison, M., Henderson, F., Jessop, A., McAteer, E., Milner, M., Neil, D., Pflücke, T., Pollock, M., Primrose, C., Richard, S., Sclater, N., Shaw, R., Tickner, S., Turner, I., van der Zwan, R., Watt, H. (1995) Using Learning Technologies: Interim Conclusions from the TILT Project. TILT University of Glasgow.

Farby, B., Land, F., Target, D. (1993), How to assess your IT investment, Oxford: Butterworth. Heinemann.

Levin, H.M. (1983) Cost-Effectiveness: A Primer Sage Publications. London.

Designing Experiments

Robson, C. (1990) Designing and Interpreting psychological experiments. In Preece, J. and Keller, L. Eds. Human Computer Interaction. Prentice Hall. Herefordshire. UK

Ethnography

Fetterman, D., M. (1998). Ethnography Step by Step, 2nd Edition. London: Sage.

Fetterman, D., M. (ed) (1984). Ethnography in Educational Evaluation. Beverly Hills, CA: Sage.

Hammersley, M., and Atkinson, P. (1983) Ethnography Principles In Practice. London: Tavistock.

Hammersley, M. (1992) What's Wrong With Ethnography? London: Routledge.

Hughes, J.A., Sommerville, I., Randall, D. and Bentley, R. (1993). Designing with ethnography making the work visible. Interacting with Computers, 5, 2, p239 - 253.

Jordan, S. and Yeomans, D. (1995). Critical Ethnography: problems in contemporary theory and practice. British Journal of Sociology in Education, 16, 3 p 389 - 408.

Zaharlick, A. (1992). Ethnography in Anthropology and its Value for Education. Theory into Practice, 31, 2 p116 -125.

Focus Groups

Morgan, D. L. (1988) Focus Groups as Qualitative Research. London: Sage. ISBN: 0 8039 3209 X

Interviews

Kvale, S. (1996) InterViews: an introduction to qualitative research interviewing. London: Sage Publications ISBN 0 8039 5820 X (A very rich book, gives a lot of food for thought, with some good guidelines)

Lockwood, F. (1991) Data collection in distance education research: the use of self-recorded audiotape. Open University H801 'Foundations of Open and Distance Education' Block 4: research and evaluation methods in open and distance education. from their MA in Open and Distance Education course, Institute of Educational Technology, Milton Keynes. (Describes a study in which he used three kinds of interview with the same list of questions. Interesting)

Patton, M. (1987) How to use qualitative methods in evaluation. London: Sage. ISBN 0 8039 3129 8 (Quite old, simple and surprisingly useful - summary of tips at the back.)

Rubin, H.J. and Rubin, I.S. (1995) *Qualitative Interviewing: The art of hearing data*. London.. Sage. ISBN 0-8039-5095-0

Questionnaires

Carter, M.P. and Williamson, D. (1996) *Questionnaire Design*. Staffordshire University Business School, Leek Road, Stoke-on-Trent ST4 2DF, United Kingdom http://www.staffs.ac.uk/buss/bscal/mandev/m_qm/t_que/que.htm

Gilbert, N. (1993) *Researching Social Life*. Sage Publications, London.

Kirakowski, J. (1997) *Questionnaires in Usability Engineering. A List of Frequently Asked Questions*, Human Factors Research Group, Cork, Ireland. <http://www.ucc.ie/hfrg/resources/qfaq1.html>

Moser, C.A. and Kalton, G. (1979) *Survey Methods in Social Investigation*. Gower Publishing Company, Aldershot, England.

Oppenheim, A.N. (1992) *Questionnaire design, interviewing and attitude measurement*. Pinter, London.

Wilson, N. and McClean, S. (1994) *Questionnaire Design: A Practical Introduction*. University of Ulster. Copies available from: UCoSDA, Level Six, University House, Univ. of Sheffield, Sheffield S10 2TN, Tel: 0114 272 5248, j.burgan@sheffield.ac.uk

Resource Questionnaires

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, p103 - 113. This includes information about a complete resource questionnaire.

Split Screen Video

Davidson, K. & Goldfinch, J. (1998) *How to Add VALUE in Moge*, N. (ed) LTDI Evaluation Studies, ISBN 0 9528731 5 X.

Jorgensen, A.H. (1990) Thinking aloud in user interface design : a method promoting cognitive ergonomics, *Ergonomics*, 33, 4, p501 - 507

Wright, P.C. and Monk, A.F. (1991) The use of think aloud evaluation methods in design, *SIGCHI Bulletin*, 23, 1, p 55-57.

System Log Data

Fritze, P. and McNaught C. (1996) *VisMap: A CD for Exploring Evaluation Issues in Computer-Facilitated Learning*, Paper presented at From virtual to reality. Apple University Consortium Academic Conference, The University of Queensland, 24-27 September 1996, published on CD-ROM and at <http://www.uow.edu.au/auc>

Gunn, C. and Brussino, G. (1997). *An Evolutionary Approach to CAL*. Active Learning, 6.

Planning Trials

Preece, J., Rogers, Y. Sharp, H., Benyon, D., Holland, S. and Carey, T.(1994) *Human-Computer Interaction*. Addison-Wesley. UK

Statistics explained simply

Barnett, V. (1991) *Sample Survey principles and methods*. Hodder publisher. ISBN: 0 340545534

Transcribing Evaluation Data

Seale, C. (ed) (1998) *Researching Society and Culture* Silverman's chapter provides a guide to transcription conventions. London. Sage publications.

Edwards, J.A. and Lampert, M.D. (eds) (1993) *Talking data: transcription and coding in discourse research*. Hillsdale, N.J. Lawrence Erlbaum. Information relating to discourse analysis conventions.

Writing Evaluation Reports

Morris, L.L. (1987) "How to Communicate Evaluation Findings." Sage Publications, London

Power On: New tools for teaching and learning Congress of the United States, Office of technology assessment (1988)

Gayle Calverley
g.j.calverley@acs.hull.ac.uk
Academic Services Learning Development
The University of Hull. HU6 7RX
236-250 Holloway Road. London, N7 6PP.
tel: 01482 465334
fax: 01482 466340

Grainne Conole
g.conole@unl.ac.uk
Learning Technology Advisor,
LaTID, The Learning Centre
University of North London,
tel: 0171 753 3109
fax: 0171 753 5012

Philip Crompton
pc4@stir.ac.uk
Research Fellow, Institute of Education
Pathfoot Building A24, University of Stirling
Stirling
tel: 01786 467615
fax: 01786 467633

Kirsty Davidson
k.davidson@napier.ac.uk
Sighthill Campus
Napier University
Edinburgh EH11 4BN
tel: 0131 455 3509
fax: 0131 455 3485

Judith George
J.W.George@open.ac.uk
Depute director, Scottish Region,
Open University, 10 Drumsheugh Gardens
Edinburgh. EH3 7QJ
tel: 0131 226 3851
fax: 0131 220 6730

Judy Goldfinch
j.goldfinch@napier.ac.uk
Senior Lecturer Dept of Mathematics,
Craiglockhart Campus, Napier University
219 Colinton Road,
tel: 0131 455 4630
fax: 0131 455 4232

Cathy Gunn
ca.gunn@auckland.ac.nz
Education Technologies Adviser,
CPD, University of Auckland
New Zealand
tel: 64 9 373 7599 ext 8354
fax: 64 9 373 7474

Chris Jones
C.R.Jones1@livjm.ac.uk
Dept of Politics, 15 - 21 Webster St
Henry Cotton Building
Liverpool John Moores University,
Liverpool. L3 2ET
tel: 0151 733 6183

Liz Leonard
E.Leonard@bio.gla.ac.uk
TLS, University of Glasgow
69 Oakfield Avenue
Glasgow G12 8LS
tel: 0141 330 4997
fax: 0141 330 4987

Ann Lewins
a.lewins@soc.surrey.ac.uk
CAQDAS Networking Project
Dept of Sociology, University of Surrey
Guildford GU2 5XH
tel: 01483 259455
fax: 01483 259551

Gaye Manwaring
g.manwaring@norcol.ac.uk
Northern College, Gardyne Road
Dundee DD5 1NY
tel: 01382 464 360
fax: 01382 464 900

Patrick McAndrew
patrick@icbl.hw.ac.uk
Manager, Institute for Computer Based Learning
Heriot-Watt University,
Edinburgh EH14 4AS
tel: 0131 451 3286
fax: 0131 451 3283

Erica McAteer
E.McAteer@udcf.gla.ac.uk
TLS, University of Glasgow
69 Oakfield Avenue
Glasgow G12 8LS
tel: 0141 330 4997
fax: 0141 330 4987

Hamish Macleod
H.A.Macleod@ed.ac.uk
Senior Lecturer, Department of Psychology
University of Edinburgh
Edinburgh. EH8 9JZ
tel: 0131 650 3444
fax: 0131 650 3461

Colin Milligan
colin@icbl.hw.ac.uk
TALISMAN
Institute for Computer Based Learning
Heriot-Watt University,
Edinburgh. EH14 4AS
tel: 0131 451 3280
fax: 0131 451 3283

John Milne
j.d.milne@aberdeen.ac.uk Centre for CBL in Land
Use and Environmental Sciences (CLUES),
MacRobert Building, Aberdeen University
Aberdeen, AB24 5UA
tel: 01224 273766
fax: 01224 273752

Nora Moge
nora@icbl.hw.ac.uk
LTDI Co-ordinator
Learning Technology Dissemination Initiative
ICBL, Heriot-Watt University
Edinburgh, EH14 4AS
tel: 0131 451 3280
fax: 0131 451 3283

Robert Newton,
r.newton@rgu.ac.uk
Dept. of Information and Media
Garthdee Road
The Robert Gordon University
Aberdeen. AB10 7QE
tel: 01224 263900
fax: 01224 263939

Martin Oliver
m.oliver@unl.ac.uk
Research Fellow, LaTID, The Learning Centre
University of North London,
236-250 Holloway Road
London. N7 6PP.
tel: 0171 753 3109
fax: 0171 753 5012

Robin Shaw
r.shaw@elec.gla.ac.uk
TLTSN Consultant
Robert Clark Centre for Technological Education
University of Glasgow
66 Oakfield Avenue, Glasgow G12 8LS
tel: 0141 330 4844
fax: 0141 330 4832

Helyn Thornbury
Helyn@mansci.strath.ac.uk
Dept of Management Science,
University of Strathclyde
40 George Street,
Glasgow G1 1QE
tel: 0141 548 4543
fax: 0141 552 6686

Gwen van der Velden
G.M.Vandervelden@ukc.ac.uk
Teaching & Learning Technology Support Network
TALENT project coordinator (TLTP3 project)
Education Support Services, Computing Lab
University of Kent, Canterbury
tel: 01227 827619
fax: 01227 762811