Building a Rigorous Research Agenda into Changes to Teaching

Mats Daniels[1], Marian Petre[2] and Anders Berglund[1]

[1]Department of Computer Systems, Uppsala University, Sweden
[2]Centre for Informatics Education Research, Faculty of Mathematics and Computing, Open University, U.K.

Abstract:

Building research into teaching has particular value in the current academic climate; this paper considers why and reflects on the particular need for Computer Science education research in its own right. Co-operation between research and teaching is needed in order to understand learning in the Computer Science context. The accelerating convergence of technologies for computing, communications and teaching affords an opportunity to integrate research and teaching objectives in Computer Science education. The paper presents four example projects in which research was tied to changes in teaching, so that general lessons could be drawn from individual experiences.

1. Introduction: Why integrate research into CS teaching?

In many Computer Science (CS) departments, there is a tension — often divisive — between teaching demands and research expectations, and a concomitant failure of communication between research and teaching. Yet understanding learning in context cannot be done without co-operation between the 'two halves' (research and teaching), in two forms: technical research (about the content of the discipline) and educational research (about the learning of the discipline). Technical research has an impact on teaching objectives, most effectively if we understand the learning process — through CS educational research — in order to bridge the gap between research outcomes and input to teaching. Making that bridge can improve motivation among 'research-oriented' teachers and increase 'clarity among teaching-oriented teachers. Computer Science educators often claim their teaching to be 'research-led', but they typically mean led by CS technical research; by ignoring CS educational research, we impoverish our provision. Computer Science is sufficiently distinctive as a discipline to require that educational research come from 'within' the discipline. The accelerating convergence of technologies for computing, communications and teaching affords an opportunity to *integrate* research and teaching objectives. This paper presents the case for doing so, and then presents some example projects in which this is being done.

The effort to marshal the latest technology and to implement the latest teaching methods to serve educational aims must be balanced with the need to seek out and address questions about which concepts, strategies, and techniques are fundamental to a Computer Science education. The fast pace of technological change poses a double challenge for Computer Science education: developments affect both the subject and the mechanisms of teaching. Educational methods race to keep pace with the opportunities afforded by technology. We must understand 'what Computing is' in order to teach it — we must marshal appropriate tools and methods to teach it well — and what we teach will influence what Computer Science becomes. This requires that research looks deeper than merely evaluating implementations, deep enough to examine what changes in teaching practice reveal about underlying issues such as concept acquisition, development of skills and expertise, sources of misconception and superstition, learning processes, the roles of different types of interaction between teachers, students, and materials, and so on. We need to know not just the effect of introducing new technology or methodology, but also the price.

1.1 What distinguishes the current academic climate?

As computers infiltrate virtually every domain, the demand for education in computing concepts and skills increases. Computer Science education has knock-on effects to all other domains, not just to engineering domains. The past several years have seen changes in the academic climate in many countries, characterized roughly by:

- higher demand but less funding, and hence pressure to find, attract, and take on ever more students, coupled with the difficulty of re-adjusting the distribution of funds between expanding and declining areas;
- the ambiguous status of teaching: whether to provide education or to provide training for work; whether to set standards in a discipline of thinking or to satisfy student 'customers';
- changing modes of study: more re-education, more mature students, more non-majors, more hybrid degrees and study programmes.

Often, the university administration looks to technology as a panacea; the scramble to offer 'distance education' courses via the Web is an obvious example.

1.2 What distinguishes teaching in CS?

In Computer Science, the fast pace of change is not just technological, but also intellectual and methodological. The discipline of Computer Science, without a firm traditional underpinning or a firm educational tradition, is buffeted by changing definitions of the domain itself. The academic discipline is characterized by many tensions: between science and engineering; between theory and practice; between training and education. The tensions are exacerbated by the current climate; in the face of income-oriented institutional perspectives, the push to satisfy future employers, the competition for students, and so on, the tensions are a matter of continual debate. Hence, the discipline is characterized by an almost unmanageable diversity:

- *academic perspectives:* Degrees and courses in Computer Science cover a wide range of goals and values.
- *representation systems:* Changes in notations and programming paradigms are attended by the need to comprehend and have competence with more than one.
- *technical context:* Our artefacts must be understood in the current technological context.

This diversity and the pace of change mean that, not only must we provide students with a solid foundation, but we must also equip them for continual learning subsequently.

More fundamentally, the nature of what we study — that our tools are also our objects of study and are also a means of teaching — sets Computer Science apart:

- The objects of study, our artefacts, are abstract and difficult to observe. They are different from physical artefacts (e.g., what does a compiler look like? or an operating system?), although they interact with objects in the physical world.
- Concepts and artefacts encompass many levels of abstraction and complex inter-relations.
- Our artefacts are dynamic; we must reason not just about their properties, but about their behaviour (potentially complex behaviour) in time.

In this context of abstract, difficult to observe, dynamic, interacting objects of study, it is a particular challenge for educators to make theory concrete — without confusing technology with theory.

More fundamentally still, Computer Science is about thinking. The constraints to thinking within the discipline are not physical, but human: our artefacts are constrained primarily by our ability to invent. Hence Computer Science teaching is about what computer scientists have managed to think about so far, and in what manner: algorithms, paradigms, languages, engines, tools, solutions are all thought products. But they are thought products that interact crucially with the physical world, and the relationship between the reasoning discipline of Computer Science and its technology is central to its particular character.

1.3 Why research CS Ed?

Even so, why should Computer Science educators involve themselves in Computer Science education research? It makes sense for them as members of a discipline: both to inform their practice, and to draw upon their practice to inform discourse in their discipline. Even though not all CS educators need *engage* in CS education research, all teachers should be *aware* of the research that pertains to their practice — as part of professional "scholarship".

Given the current conditions, it is especially important to distinguish truth from assumption, to have practice that is well-founded. Evolving teaching practice is normal to good teaching, but evaluation reliant on anecdote is not good enough. Adding a research perspective allows educators to learn more from their practice, e.g., to consider how much of local practice generalizes, and to identify the important parameters governing effectiveness in given situations. It allows educators to combine individual experiences in a meaningful way in order to address bigger issues, e.g., assessing the balance — understanding the trade-off — between practice and theory in courses and programmes.

Moreover, undertaking CS education research makes sense for many educators as individuals, adding a dimension to their practice of teaching. It focusses on questions close to their vocation, and close to their daily work. It allows their investment in teaching to contribute to their research record, one of the principal factors by which they are judged.

1.4 Why build research into teaching?

Combining research and teaching objectives has a number of advantages: economy of effort, opportunity, credibility, and adding value.

- *economy of effort:* Building a research agenda into changes in teaching combines two sets of objectives in one activity. Education's greatest research asset is its students and teachers; moreover, teaching practice is a readily accessible research resource.

- *credibility:* The evidence which arises from integrated research can provide solid ground in discussions about changes to courses and programmes. Practice-based research can contribute to arguments for support by providing credible evidence of improvement in outcomes. Well-designed studies are better than anecdote because they are documented and repeatable.
- *opportunity:* Although less controllable, the classroom has some advantages over an education laboratory. It affords the potential for discovering subtle interactions (e.g., contributing factors which are not visible without some change to practice or without *in situ* observation) and effects over time. Combining objectives brings the opportunity to build the bridges between practice and theory (and theory and practice), and to relate theory to CS education in particular.
- *added value:* Examining educational practice with a 'research eye' and a researcher's toolkit can add value to changes in practice by providing the quality of documentation necessary for a reliable transfer of experience, providing the basis for establishing the generalizeability of experience, exposing students to research. Hence, educators can see greater outcomes of local changes and efforts.

2. Rigour? — what and how?

The key question is how to achieve rigour in the face of human complexity and variability, and subject to the practical constraints of the educational setting. Rigour is plausible when research is viewed as a means of *learning* (i.e., adding information to the discourse in the community), rather than a means of *proving*. Scientific research is often not so much a process of getting answers as one of finding even better questions. This view leads to a healthy pragmatism, based on identifying the research question, considering what evidence is sufficient to address it, and accepting that constraints (and consequently trade-offs) are inevitable in the asking. It admits a variety of methods, theories, and accounts.

2.1 Accommodating different perspectives; combining techniques

There may be no single reality to which claims made in research reports correspond;
phenomena are 'constructed' not just 'discovered'. Rarely can one sort of evidence reveal everything that an educator wishes to know about the impact of introducing a change to practice, but less-than-exhaustive evidence may provide what an educator *needs* to know in order to gain insight about educational roles and value. A single technique reveals something of what's needed, and a broader examination often relies on a succession of techniques that build a collection of evidence. Research can take place in a variety of formal and informal settings, including classroom and laboratory. It can also be conducted according to a variety of learning, teaching, and methodological paradigms.

The space of research techniques is large (for a detailed exposition, see, e.g., [6]; [5]; [7]), including:

- case studies
- observation (natural tasks)
- think-aloud tasks
- in-depth interviews
- analysis of 'naturally occurring' artefacts such as documents and designs
- questionnaires & surveys
- diary studies
- focus groups
- walk-throughs
- constrained tasks
- quasi-experiments
- laboratory experiments

Overcoming the constraints of the educational setting is often a matter of combination: accumulating results from a series of studies; compiling results from studies at different sites; mixing qualitative and quantitative techniques (e.g., [3] and [2]). Combining methods allows a sort of 'triangulation' among multiple perspectives, which can achieve a more complete account. Consolidating results from different sources can improve representativeness and strengthen conclusions, as long as the data arises from cognate studies.

2.2 What gives research value?

Typical criteria for research are relevance, importance or significance of topic, and contribution to existing knowledge. The value of a research contribution rests on its validity, the extent to which an account accurately represents the phenomena to which it refers. Our aim must be to establish that research results are valid 'beyond reasonable doubt': that they are plausible, credible, and supported by well-documented evidence. The value of evidence relies on clear and honest reporting of data that has been collected in a systematic manner using appropriate qualitative and/or quantitative methodologies. The crucial issues for research are whether it is representative, generalizeable, replicable, predictive, honest — and whether it recognizes its own assumptions and considers alternative interpretations or accounts. 'Good' evidence is appropriate evidence, i.e., data relevant to the question; 'bad' evidence fails to provide relevant information. We learn from 'failure' (i.e. refutation of our hypotheses) if the study is well-reported. Flawed evidence (i.e., evidence which is compromised by some limitation of the study) can still be useful, as long as we understand its limitations. A pragmatic approach to rigour relies on researchers knowing the value of different sorts of evidence and using available evidence within its value.

3. Examples

The section presents four example projects in which research is being tied to changes in teaching, so that general lessons can be drawn from individual experiences.

3.1 Example 1: The Runestone Project

The Runestone project [4] is a co-operation between Uppsala University, Sweden, and Grand Valley State University, Allendale, Michigan, USA. The basic idea is to develop and evaluate the notion of incorporating international group projects into the undergraduate Computer Science curriculum. New dimensions to student teamwork are added, requiring students to handle collaboration that is remote (distance and time), cross-cultural, and linguistically challenging. Runestone runs for three years, with the prototype version running in winter 1998. The 1998 pilot study will be followed by a full-scale implementation in 1999 and another in 2000.

The pilot study consists of four Swedes working with four Americans on a group project that involves controlling a small steel ball on a tilting board that is physically located in Sweden. The work started in the first part of January and will be finished by the end of March. For the Swedish students, it is part of a course that started in September, whereas for the Americans it is the major part of a course that started early January.

Runestone's aims include:

1. Giving our students international contacts and experience with team-work with people from a foreign culture and with a different educational background.
2. Increasing peer-learning.
3. Benefiting staff by close collaboration with other universities, giving insights to other departments and ideas for new teaching methods.
4. Giving experience with use of new techniques in the running of a course, both for teachers and students.

The project encompasses both evaluation and research.

3.1.1 Evaluation:

The evaluation component will both assess Runestone in terms of its own aims, and draw from the experience any lessons that enable the Runestone format to be transferred to other departments and institutions. The evaluation will for this reason aim to distinguish between domain-specific and general lessons. Key evaluation issues include: the efficacy of the method — whether students learned and to what extent they collaborated; the cost of using this form of education, both in money and time, both for students and staff, e.g., how much time is spent on becoming acquainted with new techniques for communication; whether (and if so how much) language problems impede progress; which aspects of the set-up contribute to or detract from the collaboration.

3.1.2 Research:

The two major aspects of the research component are peer learning and (changes in) cultural attitudes. The research will use largely qualitative methods to examine the processes students use in pursuing the project, the patterns of communication among the students and what their communication reveals about their attitudes and understanding.

One hypothesis of the project is that the different educational backgrounds between the two sets of students will promote interaction, because each group will need to draw on the other's knowledge. Another hypothesis is that peer-learning will contribute substantially to students' understanding during the project. There are likely to be plenty of occasions for these students to explain things to each other. The project will examine occurrences of peer learning in order to try to determine what characterizes effective occasions (i.e., those in which understanding is improved), whether there are patterns of occurrence, and what factors might be used to promote effective peer-learning. The anticipation is that peer-learning will occur both spontaneously and in more formal settings when students in one group teach students in the other about topics not covered at that site.

There are several sources for data collection in this pilot study.

- There are two weekly meetings that are audio recorded and followed up by a questionnaire. The questionnaires are one-page and focus on that meeting, asking about the organization of the meeting, the outcomes (decisions, learning, conflict resolution, clarification, etc), and the students' satisfaction with the proceedings, both overall and in terms of their own role in the meeting. The first of these meetings is with the students at both sides of the Atlantic, both ends taped. The second meeting is with the teachers during which students and teachers reflect on the course of the project during the week. This 'de-briefing' is loosely structured, following a standard script, but allowing the teachers to respond to developments or observations.
- The students, both the ones involved in the pilot study and the other students at the Swedish side fill in weekly project logs which ask for a daily log of their time on the project, their activities and interactions during that time, and the outcomes.
- The teachers are asked to keep a diary of their observations, making particular note of any evidence of peer learning, culture clashes or development of sensitivity, collaboration, effective or ineffective procedures, and technology issues.
- All student mail relating to the Runestone project and IRC logs from the weekly student meetings are collected and copied to a team member in the UK, who will keep an eye on it and ultimately analyse it for

things like decision strands, student roles, evidence of peer learning, and culture issues.
- Video tapes of the initial and final video conferences, as well as entry and exit questionnaires completed by all students.

The analysis of the pilot study will be qualitative and largely data-driven within the key topics. Results from the pilot study will be used to design more focussed research instruments for the subsequent years, facilitating more efficient analysis on a much larger student pool.

3.2 Example 2: Changing assessment methods in order to influence study habits

Our experience shows, and there is much corroborative evidence from research, that students tend to adapt to the requirements of the educational system. That is, students do what they need to do to pass the examinations and ultimately to obtain a diploma. This is undoubtedly a rational behaviour from the students' point of view, but it has some clearly undesirable effects, like a tendency among these students to study hard before the examination, to learn to solve 'typical problems', and to prefer learning facts over trying understanding the underlying ideas in the subject they are studying.

In a project that is carried out at the Department of Computer Systems at Uppsala University, Sweden [1], we are trying to break these patterns, by changing the examination forms. By changing the assessment methods, we hope to change students' study habits (e.g., encouraging them to study during the whole course, rather than mainly right before the exams) and hence to promote deeper learning.

The assessment methods being tried in this project are:
- weekly assignments instead of a final examination
- seminars with a public debate on smaller projects
- examinations with questions that are focusing on the higher levels in Bloom's hierarchy.

The methods are used on third-year students taking 'Algorithms and Data Structures', 'Computer Architecture', and 'Signals and Systems' in the Engineering Physics program.

In order to evaluate the effects of the assessment methods, and their influence on the students' study habits, data is collected through questionnaires and interviews with students who are selected to be a 'representative' sample of the students in the courses concerned. There are group interviews, focusing on the students' opinions of the methods used, as well as individual interviews, in which the emphasis is on students' understanding of fundamental concepts in the courses. Students are asked to keep 'study logs' which record how much time they spend on particular activities for the different courses. Since we are changing the assessment methods (and since the assessment method in a sense 'produces' the grades), comparing grades with earlier years (when using another method to 'produce' the grades) gives very little evidence about the students' understanding. Thus, other methods than simple comparison of grades must be used.

If we manage to show that the students' learning is good and that their motivation and study strategies are better, then we clearly have interesting results to use, to communicate, and to propose to others to adopt. If there is no change to be seen, we have an interesting piece of evidence to discuss on the issue of examination. However, preliminary data clearly indicate that the students' habits are improving, that is, that the methods, in some sense, improves the education.

3.3 Example 3: Entry-level course Internet presentation trials

The Open University (OU), which teaches around 150,000 students at a distance, has developed a well-tuned machine for providing high-quality university education for part-time students studying at a distance. The Internet trials [8], [10], [9], and [11] examined whether the Internet could be used effectively (including cost-effectively) as the delivery medium for that education. For two years, members of the Computing Department have been developing a learning environment to support the whole instruction process, encompassing students, teachers, staff support, and administration. We have investigated mechanisms for:

- interactions among students and tutors via email, conferences and Web resources;
- assignment marking using an electronic marking tool (with a component for quality monitoring);
- an electronic assignment handling system, including electronic assignment submission, and automatic verification and record-keeping;
- synchronous and asynchronous Internet-based problem sessions;
- a Web-based, automatic registration system (see http://mzx.open.ac.uk); and
- electronic examination using encrypted examination papers downloaded via the Web at strictly supervised examination centres at appointed times.

The systems have been tried on an entry-level and an upper-level Computing course, involving approximately 350 students and 23 experienced tutors in 1996, and some 500 students in 1997.

Coupled with the Internet presentation trials were both *evaluation* (we wished to gather enough data to make well-founded comparisons between conventional and electronic delivery, assessing for example the appropriateness of the technology, comparative costs, and the tutors' and students' experience of the course) and *research* (addressing more general questions about the impact of the medium and the nature of interactions between students and teachers, investigating concept acquisition, student differences, the impact which individual differences in tutor style had on resource usage and on learning effects, which changes in culture help to preserve or improve teaching quality while adapting to screen-based and often asynchronous interactions).

Considerable data, both qualitative and quantitative, was collected to support well-founded comparisons of conventional and electronic delivery:

- all examination and assessment results
- about 3,000 marked assignments
- questionnaires from both students and tutors, covering background, experience, prior education, attitudes, and their experience of the course
- mid-term and end-of-term exercises by students to assess skills and concept acquisition
- all conference postings and most electronic mail
- interaction logs completed by the tutors (contemporary records of their interactions with students)
- tutorial logs completed by the tutors (describing their conduct of problem and discussion sessions and their observations about what happened)
- records of de-briefing meetings with tutors

Data analysis used various methods, both quantitative and qualitative, to address different issues, such as the nature and level of feedback on assignments; patterns of communication in electronic mail; and models for electronic problem sessions, as well as factors influencing effectiveness.

The results of the many strands of analysis can be summarized with observed costs and gains:

3.3.1 Costs

- Substantially more technical support is required.
- Tutors bear the brunt of the transition: mastering new tools and skills, evolving a new culture, new strategies, new materials.
- Students must take responsibility for their own learning. Some presentation costs (e.g., connect time; printing) are off-loaded onto students.
- Students were disappointed in their interactions with other students; with limited resources, this is a difficult medium in which to establish a 'community of learning'.
- Electronic tutorials are as yet no substitute for face-to-face interaction, although they clearly have value and tremendous potential.

3.3.2 Gains

- More rapid feedback for students.
- Increased tutor collaboration and communication.
- Greater access for students.
- Increased administrative efficiency.
- Reduction in administrative errors.
- Potential for flexibility.

Supported Internet presentation is not a cheap option, but it may be one that can provide greater flexibility and can shift effort from administrative details to teaching.

3.4 Example 4: AESOP: An electronic student observatory project

AESOP ('An Electronic Student Observatory Project') [12] is a collection of computer-based data collection tools for instruction and research. This project aims to use research instrumentation integrated into the learning environment provided to students to investigate students' learning of computing concepts and programming skills as well as their use of resources.

Our educational environment is one in which over 5,000 students study independently at a distance, off-line, using software developed for an entry-level, distance education course in Computing. The course aspires to innovate delivery, incorporating object-oriented teaching in a Smalltalk environment, the option for fully electronic presentation, a group project requiring students to co-operate in producing software using computer-mediated communication, automatic assignment submission and return using automatic processing, and network access to teachers. Students and teachers interact primarily through e-mail.

We wish to observe these students unobtrusively, electronically, and automatically, and to record the observations in a manner that is useful for both instruction and research. The computer-based learning environment is being instrumented to provide automatic collection and analysis of data about student behaviour, from logs of network access to traces of software development and modification. Our recorder must create a transcript short enough to be sent via e-mail (by students who pay phone charges), readable both by humans and automated analysis tools, and replayable, so that the student's session can be reproduced on an observer's computer.

AESOP's recorder and replayer can be used as instructional tools, allowing a student to record a session in which some difficulty arises and send it to the tutor with annotations that appear when the tutor replays the session. The tutor can record an alternative 'idealized' session with annotations, which the student then replays. Hence, student and tutor can develop a remote dialog using the actual programs and learning environment.

AESOP's tools can be used for research. Significant events, such as errors, can be identified as they happen, observed in the transcript, and searched for by eye or by using the analysis 'toolkit' (still under development). This will enable researchers to search automatically through multiple transcripts and obtain statistics about interesting events. The analysis 'toolkit' will incorporate filters, profilers, and compilers for the data, in order to minimize the path between posing a question and finding the evidence available in the data. In this way, we hope to identify emergent patterns which will provide insight into how students learn to program. For example, we hope to re-construct the processes students use in creating and debugging programs, to correlate interaction behaviour to group performance, to investigate concept acquisition and change, to analyse factors affecting performance, to identify points of difficulty or misconception, and so on.

Most studies of Computing education are based on small groups (for example, the members of one or two undergraduate classes, usually not more than 25 subjects per group); many are largely qualitative, collecting limited quantitative data from pre- and post-tests only; most are based on observations in group settings (e.g., the problem session) or in artificial settings (e.g., laboratory tasks); very few are longitudinal, and many rely on a single task. The electronic observatory will allow quantitative data collection on an unprecedented scale, potentially covering the normal, independent behaviour of thousands of students per year, over the life of the course.

4. Summary

Where there is challenge, there is also opportunity, and the convergence of technologies is an opportunity to integrate Computer Science research and teaching objectives. This paper has presented examples of projects which attempt such an integration. These attempts are not focussed on how to use new technologies, rather on understanding the learning process within Computer Science education, often in the light of introducing new technologies. The application of new technologies can provide a natural vehicle for gathering data for Computer Science Education research.

Effective research in CS education requires knowledge of the subject, Computer Science, as well as knowledge of the teaching and learning processes. The goals of the sort of 'embedded' CS education research described here are 'close to heart' of CS educators, and perhaps the most immediate gain is an increased insight for the researchers themselves into the teaching process of the subject. But by building rigorous research into teaching, the insight has value beyond the particular situation:

- results are generalizeable,
- results can be compared to results of other studies,
- the work is documented and repeatable, and
- the results have a credibility far in excess of individual anecdote.

These benefits are achieved using the ”normal” teaching situation as vehicle.

Acknowledgements

Obviously, none of the projects mentioned is the work of one person. The authors acknowledge the project teams:

- Runestone team: Carl Erickson, Lars Asplund, Mary Last, Bruce Klein, Christina Björkman, Vicki Almstrum.
- Assessment methods team: Shirley Booth, Anders Svärdström, Christina Björkman, Annika Lundmark, Vicki Almstrum, Vigner Karlsson.
- MZX Internet team: Pete Thomas, Blaine Price, Linda Carswell, Mike Richards.

- AESOP team: Pete Thomas is the project leader; Mike Martin and Malcolm MacGregor are collaborators, as are Mark Woodman and the M206 course team.

The authors are particularly grateful to the students and teachers who have participated in the various studies. Thanks also to Sally Fincher.

References

[1] Berglund, A., and Daniels, M. (1997) Improving education quality, a full-scale study. In: Proceedings of ACM SIGCSE Symposium (San Jose, February).

[2] Brannen, J. (Ed.) (1992) Mixing Methods: Qualitative and Quantitative Research Theory and Practice. Avebury.

[3] Bryman, A. (1988) Quantity and Quality in Social Research. Allen and Unwin.

[4] Daniels, M., Almstrum, V., Asplund, L., Björkman, C., Erickson, C., Klein, B., and Petre, M. (in press) Runestone: international student collaboration on projects. To be presented to: IEEE Frontiers in Education Conference (Tempe, Arizona, November 1998).

[5] Denzin, N.K., and Lincoln, Y.S. (Eds.) (1994) Handbook of Qualitative Research. Sage.

[6] Isaac, S., and Michael, W.B. (1989) Handbook in Research and Evaluation for Education and the Behavioural Sciences. EdITS Publishers.

[7] Gall, M.D., Borg, W.R., and Gall, J.P. (1996) *Educational Research: An Introduction*. (Sixth edition) Longman.

[8] Petre, M., Carswell, L., Price, B., and Thomas, P. (1997) Innovations in large-scale supported distance teaching for the Internet: transformation, not just translation. In: Proceedings of IEEE Frontiers in Education Conference (Pittsburgh, November).

[9] Petre, M., and Price, B. (1997) Programming practical work and problem sessions via the Internet. In proceedings of: ACM SIGCSE/SIGCUE Conference on Introducing Technology into Computer Science Education (ITiCSE '97), Uppsala, June. *SIGCSE Bulletin*.

[10] Price, B., and Petre, M. (1997) Large-scale interactive teaching via the Internet: experience with problem sessions and practical work in university courses. In proceedings of: ED-MEDIA/ED-TELECOM '97 Conference (Calgary, Canada, June).

[11] Thomas, P., Carswell, L., Price, B., and Petre, M. (in press) A holistic approach to supporting distance learning using the Internet: transformation, not translation. *British Journal of Educational Technology* (first of a 4-part series).

[12] Thomas, P., Martin, M., MacGregor, M. (in press) AESOP: An Electronic Student Observatory Project. To be presented to: IEEE Frontiers in Education Conference (Tempe, Arizona, November 1998).