Introduction

Learning to be a researcher is a process of apprenticeship, the student working alongside an experienced practitioner and picking up the trade by example and personal guidance. This certainly seems to be the most expeditious way to develop an understanding of the subject matter in a chosen research area, as well as to assimilate some of the more general points of the researcher's craft — a knowledge of the literature, what's worth reading and what isn't, a nose for worthwhile problems, good research taste. Half the result is uneven, many essential skills are not even perceived as necessary, let alone imbibed, until researchers are well into their career. Is it possible to teach research skills to a graduate class in computer science? What should be taught, and how?

This article describes a course that has been given for three years to incoming graduate students in the Computer Science Department, University of Calgary. The course was designed to meet a perceived need to accelerate students' progress at the beginning of their MSc program. The official course description reads:

> An introduction to and survey of research areas and methods in Computer Science. Professional skills in computer science research such as review-writing, critical evaluation, and the preparation of research proposals.

More specifically, the aims of the course are threefold:

- *To provide an introduction to and survey of research being undertaken in the Department* by faculty and other graduate students.
- *To develop students' professional skills* by providing an opportunity to practice giving research presentations, writing reviews and proposals, critical listening and note taking, and the use of document preparation facilities for thesis preparation.
- *To accelerate students' progress in the graduate program* by providing an opportunity to make an early start on thesis research.

These aims are discussed with students at the beginning of the course. This article concentrates on the last two.

The first component, a survey of research areas, introduces the research being pursued in our Department. It is not really designed to provide a balanced survey of the subject as a whole, but rather to perpetuate an awareness of what is going on in the Department. This component is

intended to help students identify prospective supervisors, to broaden students' horizons and encourage interaction among the graduate student body by giving a feeling for what other students are doing, and to provide examples of seminar presentation styles and techniques. Most of the first half of the course is dedicated to this aim. It is basically an ordinary seminar program by Faculty and other graduate students, and needs no further elaboration.

The other components are quite different in nature and attempt to impart professional skills for computer science research. A number of skills have been identified and are introduced below. Whether they are teachable in such a forum is open to debate, as is the question of whether important skills have been omitted. It is not clear that we have got this course right!—but it does provide a starting-point. Students are encouraged to develop their research skills through practice in a variety of ways. This culminates in the preparation of a thesis proposal, which allows them to practice their skills in the context of their own potential thesis topic.

It is worth reflecting on the extent to which professional research skills are subject-dependent. Would not a course on research in, say, Physics, look the same? And since traditional disciplines do not, in general, find such courses necessary, why is Computer Science different? To answer the first question, I expect that the research skills are similar (though certainly not identical), but that they are best taught in a concrete manner that involves examining particular pieces of work, and these belong firmly within the discipline addressed. As for the second, I believe that Computer Science *is* different from other subjects because of its youth and because existing paradigms and lines of enquiry are much more fluid and less firmly established. As Brooks (1978) argues, we work in a tractable medium and build from "pure thought-stuff"; our research is open, unfettered, creative—and, as a result, extremely hard to control.

Components of the course

As well as the area-specific information that can probably best be learned under the direction of a supervisor, the researcher needs a number of more general skills. These are identified and discussed below, and some indication is given about how they are taught. Part of the purpose of this article is to provide pointers to articles that can be used as background material, by either the instructor or the class as a whole. Sources that are worth reading in general, and are not targeted at any particular area identified below, include Bundy *et al.* (1986), which is addressed to PhD students in Artificial Intelligence at the University of Edinburgh but in fact is quite widely applicable, and Barzun & Graff (1977), which is an elegantly-written and readable account of research skills aimed principally at the historian and social scientist.

MAKING PRESENTATIONS

One essential role of the course is to give students experience in presenting material in front of an audience. Each student does one brief presentation and one longer one. Brief presentations are limited to 15 minutes and followed by a 5-minute question period. This is a common format for academic conferences. Most people agree that it is very difficult to say anything worthwhile in so short a time, but it can be done with experience and careful preparation, and the art is well worth practicing.

Full-length seminars are ideally 50 minutes long, followed by up to 10 minutes of questions. Again, this is a common format for major presentations and Departmental seminars. Everyone in the class gives a brief presentation before they undertake a long one, and we have class discussions on common problems to give students a chance to improve their presentation style. In fact, depending on the number of students in the class, time did not always permit 60-minute seminar slots, and sometimes two "full-length" seminars were squeezed into one 90-minute period, with 35 minutes each plus 5 for questions.

Each presentation is on a different topic. The brief presentations introduce and critically discuss a published research paper. It is as though the student were the author of the paper, presenting the work at a conference—although he or she is expected to be more dispassionate and critical than authors usually are! Students are strongly encouraged to follow up references to learn the necessary background. The full-length presentation focuses on a particular, well-focused, topic, perhaps presenting and contrasting the approach taken in a small group of (say 3 or 4) related papers. It is expected to involve substantial study, and students need to read fairly widely to gather and sift appropriate material. They gain additional credit for producing a handout for the class to accompany their seminar; but are instructed not to merely photocopy their overheads.

Students choose topics themselves or in conjunction with their research supervisor. They are encouraged to select papers that are in areas they already know something about or are interested in pursuing further. However, they are not allowed to present their own work or that of their supervisor. The principal aim of the exercise is to encourage students to develop a thorough understanding of someone else's work and undertake a critical analysis of it. Giving the rest of the class an interesting and educational experience is important, but secondary.

Most of the time, students are listening to others talk. They are expected to learn from this experience—not just about the subject matter but also about how to make a good presentation. They are encouraged to be critical, both of the talk and of the ideas presented, and to ask questions that are intended to clarify their understanding of the material presented. But they are warned not to destroy other people's confidence in themselves, and to discuss any criticism of another students' presentation style with them sensitively and in private. Above all, they are urged to consider how they can avoid making the same mistakes themselves!

We have experimented with videotaping presentations and giving students private access to their videotape. Although the class was initially very keen on the idea, it turned out not to be worthwhile and was dropped from later versions of the course. In practice, students did not make time to view their video critically and use it constructively. Perhaps the instructor should have gone through it with each student individually. However, it seemed better and simpler to dispense with videotaping and encourage feedback from other students in informal discussions outside the class.

Through numerous presentations by faculty and senior graduate students on the research they are undertaking, students are exposed to a wide variety of seminar styles; occasional class discussions identify and analyse differences between presenters' styles. This also provides an opportunity to practice skills of critical listening and note-taking—students are expected to make their own notes for each seminar. They must hand in a brief summary and critique of the ideas presented in a certain number of the seminars (each student chooses which ones to cover). Students are also expected to participate in the question period at the end of each seminar.

There are a number of sources of helpful information about giving presentations. Parberry (1988), although ostensibly oriented towards presentations in theoretical computer science, is pertinent to all areas of the discipline. Mental (1971) is a humorous account of some things to avoid. A John Cleese movie entitled "Please can we have that the right way round" is an entertaining and instructive illustration of mistakes (Cleese, 1976). Michaelson (1982) contains some useful pointers for students giving presentations. Although none of these sources contain any surprising insights, and all the material in them is straightforward and well-known even to the moderately experienced presenter, nevertheless students do benefit from having the "obvious" pointed out explicitly.

CRITICAL EVALUATION AND WRITING

A great deal of attention is paid in the course to critical evaluation and writing skills. The foundation for this is laid by requiring students to write two reviews on papers of their choice. They model their reviews on the more substantial ones in *Computing Reviews*. To focus attention on the requirements of the task, annotated examples of good and bad reviews are handed out and discussed in class.

It is stressed that the review must evaluate the work critically—a mere summary is not enough. It must bring out good and bad points, and relate the work to other research. Students are not graded on their opinions of the papers they review but on the arguments they recruit to support them, and the way in which they express them. Nevertheless, a copy of the target paper is submitted with the review to allow its accuracy to be assessed. Students are told that they will gain more credit for a balanced evaluation than by taking an extreme stand for or against the work. A great deal rests on their ability to write clearly, cogently, and concisely.

Normal professional standards are expected for the length and presentation of reviews. A suitable document template for the LATEX document preparation system (Lamport, 1986) is provided so that the submitted work will look very much like the publication format of *Computing Reviews*. Students are instructed to aim for around 500 words and under no circumstances to exceed 1000 words; they must include a word count at the end of the review. They should be in a position of having so much to say that it is a real challenge to reduce it to this size. Of course, reviews are annotated when they are graded to indicate, as completely as possible, their strengths and weaknesses—from critical content to minor typographical mistakes.

Students themselves choose what papers they review (one paper per review), and are encouraged to do so in conjunction with their research supervisor. They may not be the same as the papers covered in oral presentations. Appropriate choices can give a student a head start on thesis work. Also, careful coordination of the topics chosen for the presentations and reviews can make a student's job in the course much easier. Quite apart from the goal of honing critical evaluation and review-writing skills, students who choose papers intelligently will, by the end of the course, have read and thoroughly analyzed a number of key papers in their area of interest.

THESIS PREPARATION AND THESIS ORALS

There is some routine information about what should go into a thesis that it is useful to review in a course on research skills. This includes the formal sections into which a thesis is divided, the abstract, table of contents, list of figures, role of appendices, the difference between a reference list and a bibliography, and so on. There are a number of standard books on the subject (e.g. Campbell and Ballou, 1974), some of which may be officially sanctioned by a particular university. The level at which the thesis should be written—that is, the targetted audience—and the style in which it is written—namely formal, scientific description—are discussed. So is the fact that the title and abstract generally circulate far more widely than the body of the thesis itself, and the implications of this when choosing the title and writing the abstract.

The experience of examining graduate theses gives us information about the way in which they tend to be read that can be very useful to students. Students often do not realize that examiners are not compelled to read every word of a thesis in order to form an opinion of its value and quality, and that the sections may not be read in their natural order. Reflecting on the realities of the evaluation process helps to highlight the crucial importance of presentation and layout in giving a subjective impression of quality and attention to detail. The importance of the introduction and conclusion, and the fact that they should represent one's absolutely best writing style, are also worth noting.

Thesis orals are a nerve-racking experience for students; this is particularly so in universities in which orals are closed so that a student's first experience of an oral examination situation is his own. The experience can be made easier and more productive by familiarizing students in advance with the oral situation and discussing the university regulations that govern the conduct of orals. Moreover, some advice can be given on how to approach thesis orals. For example, the purpose of the examination is to reach a judgement rather than to inform or educate examiners. Questions are asked not because examiners want to know the answer, but because they want to see whether the candidate knows it. Unanswerable questions are sometimes used to see how the student reacts, how he or she approaches difficult problems. Fuzzy or ill-formulated questions are posed to see how the candidate goes about clarifying them, or because examiners are unable to clarify them themselves.

I have found it useful to run "mock" orals in which a student who is about to graduate is examined on his or her thesis by the class. This gives students some insight into the problems

of an examiner too: they must find time in a busy schedule to read the thesis critically, evaluate it, and come up with sensible questions through which they can lead the candidate into a discussion of pertinent aspects of his work. Moreover, the thesis oral is perhaps the only examination situation in which the examinee knows vastly more about the details of the subject under consideration than the examiners—this can place considerable stress on examiners and students should be sensitive to their predicament. The preparation of a standard examiner's report on the thesis is also a very useful exercise that allows students to practice their skills of critical analysis. By being forced to reflect on the value of another student's thesis, they gain insight into what makes a good thesis that should prove helpful when they come to write their own.

RESEARCH IDEAS AND RESEARCH PROPOSALS

It is extremely difficult to teach in a course how to do research; that is why we use the apprenticeship system of graduate study for this purpose. However, some basic ideas can be taught explicitly, and others can be conveyed by providing an opportunity for practice and giving feedback on the result. To provide a practical context for this activity, we focus on the research proposal as a vehicle for communicating information about good research skills.

An earlier paper addresses the writing of research proposals (Witten, 1990). This adopts the view that to do good research one must formulate a *question* that the work will strive to answer. This should not just be an isolated question, but one relating to a longer-term research *theme* that evolves over a substantial part of one's career. Moreover, it is a good idea to begin with not just a single question, but a few (although not too many) that differ in riskiness, and hence potential value. The researcher must be able to assess the value of these research questions himself or herself, in order to pick good ones and present them clearly.

The most critical issue in computer science is not so much the generation of research questions, for our field is young and there is much to do. Much more difficult is the *evaluation* of research ideas: it is crucial to be able to dispassionately analyze one's own ideas, weigh their strengths and weaknesses, sharpen them, and present them in the clearest possible way to permit their evaluation by others. These issues are discussed in more detail in Witten (1990); they underlie the the strong emphasis in this course on the critical evaluation of research work.

To help students begin their thesis research, a thesis proposal is due shortly before the end of the course. Students are told that this need *not* be on what they really end up doing for their thesis—although it is obviously to their advantage that it is. To start them on this job in good

time to write a cogent and polished proposal, a brief synopsis of the thesis proposal is required halfway through the course which sets out the goals of the thesis and gives some indication of the evaluation methodology that is proposed to test whether the goals have been met. The thesis proposal must include a literature survey, and students who choose their review and presentation papers wisely will already have a head start on this.

A large part of the job of writing a proposal is to come up with an idea for a thesis project. Students are judged partly on the idea itself, partly on the relevant background knowledge they demonstrate, and partly on their ability to write it up clearly, concisely, and eloquently. The final proposal is limited to 10 single-spaced typeset pages. This includes any figures, although the title page (with an abstract) and reference list are not counted in the total. As in any proposal, clarity and presentation is paramount.

There are several sources of advice about how to write up research proposals, and research projects in general. Langley's (1990) few pages of editorial are densely packed with pertinent and practical suggestions that I find extremely useful. Knuth *et al.* (1988) is a substantial collection of notes from a course at Stanford on "Mathematical writing"; it contains a great deal of information on technical writing and the effective presentation of mathematics and theoretical computer science. Sides (1984) gives a more general account of report-writing in our discipline.

THE PUBLICATION BUSINESS

A number of aspects of the academic publication business are discussed. First, there are the relative merits of submitting work to journals, conferences, and workshops, the standing of different journals and conferences and how this might be measured, citation indexes, the "cv value" of a paper, the role of books and edited collections. These are all questions on which we form our own opinions through experience, and are raised in order to accelerate the process by sharing and discussing the viewpoints of the instructor and the students in the class. The issues involved are clearly extremely sensitive to the instructor's bias, and other points of view are discussed too. Less contentious, but equally important, is the question of how one might go about deciding which journal or conference is suitable for a particular piece of work.

Next, the process of publication is outlined for journals, conferences, and books. This includes calls for papers, information for journal authors, the mechanics of preparing and submitting a paper, refereeing, the role of the editor, reasonable and unreasonable publication delays, the procedures of notification of decision, possible resubmission, acceptance, copy-editing, proof-reading, and index preparation.

The refereeing process is singled out for special attention, involving as it does questions of critical evaluation, professional ethics, and one's involvement in and obligation to the scientific community. Parberry's (1989) excellent article is used as a foundation for discussion; despite the title's emphasis on "theoretical" computer science almost all of it is applicable quite broadly within our discipline.

Another area where useful and specific information can be provided is in scholarly style and proof-reading. A useful guide is van Leunen's *Handbook for scholars* (1986), while students are also made aware of the standard reference *The Chicago manual of style* (University of Chicago Press, 1982). Furthermore, students should become familiar with some matters of elementary typography, discussed in a straightforward way by Williams (1990).

Finally, there are numerous ethical and societal issues in the scientific publication business. An interesting, controversial, and occasionally outrageous account of the publication "game" is available (Anonymous, 1987), along with accompanying comments by respected scientific figures; this can be used as a basis for class discussion. Crutcher (1991) gives a sardonic view of how scientists can advance their career with minimum effort. Intellectual honesty, including delicate questions of misrepresentation and even fraud, need to be addressed (Crawford and Stucki, 1990). This can easily be illustrated at the level of the computer science graduate student. For example, it is misleading to show sample screens that have been produced by hand with a drawing program rather than generated automatically by the student's program—unless they are clearly identified as mock-ups. It is misleading to imply that a system does things it doesn't actually do right now (e.g. future enhancements, even very simple ones). If there are problems with a thesis project (e.g. the system is incomplete or parts of it don't work), they should be noted and discussed in the dissertation; it is disastrous to be "found out" in an oral.

RESEARCHER'S TOOLS AND TECHNIQUES

During the course, students learn how to use a few important research tools. For example, they must become familiar with at least one comprehensive document preparation system; this will relieve their burden when it comes to actually writing a thesis. They must get into the habit of using tools like spelling checkers. They are encouraged to document conscientiously all references that they consult, in a uniform way, keeping full details of each source; most use a system like BIBTEX for this (Lamport, 1986).

Students acquire library skills and find out how to use the various abstracting and indexing publications of our field, principally the *ACM Guide to the Computing Literature*, *Computer Abstracts*, *Computing Reviews*, *Current Contents*, and the *Science Citation Index*. One possibility is to have a librarian come in to explain the library's computer science collection; another is to organize a scavenger hunt in the library to force students to become very familiar with the cataloging system and the holdings in our discipline.

Depending on the students' interests, other tools are occasionally introduced briefly in class, such as statistical analysis programs and symbolic algebra packages. They may be exposed to ideas such as exploratory data analysis in the style of Tukey (1977) and graphical information presentation skills such as the review by Tufte (1983) of the visual display of quantitative information. As in other parts of the course, the aim is to make students aware of the resources that are available should they be needed, rather than to provide comprehensive coverage of these areas.

Again depending on the ideas and interests of the students, some class discussions focus on particular skills that are relevant to the researcher, such as brainstorming techniques and writing academic curriculum vitae. Finally, it is worth sketching the opposing views of science as a rational (e.g. Popperian) endeavor and as a non-rational (e.g. Kuhnian; see Kuhn, 1970) process (Newton-Smith, 1981) in order to equip students to direct their research according to contemporary scientific thinking.

Organization and assessment

STRUCTURE

The class meets once a week for two semesters. The norm at our university is to meet twice a week for one semester, but it was decided that this type of material is best spread out over a longer time period. Classes are organized around weekly seminars by students in the course, other graduate students, and members of faculty. The first hour of each 90-minute class is generally a full-length seminar (which all graduate students and faculty are invited to attend); this is followed, after a few minutes' break, by a briefer presentation or discussion period.

The structure of the course naturally varies a great deal with the number of students enrolled in it. Once enrollment exceeds 16 or 17, the nature of the class changes because each student can only give one presentation rather than two. With smaller numbers, there is more opportunity

to include classes on diverse tools and techniques. Table 1 gives a typical schedule for an enrollment of eleven students. The 25 classes include:

- a welcome from the Head of Department and a presentation on facilities by members of the technical support staff;
- nine research seminars by Faculty and senior graduate students;
- eleven brief and eleven full-length presentations by students in the class;
- fourteen sessions led by the instructor covering the material described here;
- three brief guest presentations and a movie;
- a mock thesis oral examination.

There is not a great deal of time to cover the material reviewed in this article—only fourteen sessions, most of which are short ones. It seems much better to pack information into fast-paced classes than to run the risk of having to spread the material too thinly. Much of what students learn is through practicing relevant skills rather than listening to lectures on research methods.

Towards the beginning of the course, the full-length seminars are generally presented by faculty; around the middle, by senior graduate students; and in the later part, by students registered in the course. However, this structure is quite flexible, particularly in accommodating the needs and desires of senior graduate students to present and obtain feedback on their work. The briefer presentations are generally given by the instructor or by students registered in the course.

ASSIGNMENTS

Assessment is based on a number of small components which have been mentioned individually above. Table 2 shows the actual percentage weights used in the most recent version of the course. The individual assignments are spread as evenly as possible over the duration of the course, beginning with the brief presentation and reviews and ending with the mock oral and thesis proposal.

In an attempt to increase the amount of feedback given on class presentations, written comments on each one were prepared by the instructor and circulated to all students after the last

presentation.[1] Moreover, after each review was handed back, common problems and shortcomings were discussed in class. However, these attempts to maximize feedback were not entirely successful.

STUDENT ASSESSMENT

Student evaluation of the course has been generally extremely positive—apparently the existence of this course has elicited envy from graduate students in other departments! However, many students point out that aside from being given the opportunity to practice, no real attempt is made to *teach* skills of presentation, writing, and critical analysis.

A sampling of comments from students in the course include:

- students learn less from their peers' presentations, more from faculty and senior students;
- faculty and senior students should be more sensitive to the objectives of the course when designing their presentations;
- prompt and detailed feedback on student's work is essential;
- uniformity of grading is a potential problem because of bias due to the instructor's own research interests;
- course organization becomes unwieldy with more than about seven to ten students;
- the course is a lot of work, although this is ameliorated by the fact that students choose their topics themselves.

There is some disagreement amongst students on whether the course should spread over one term (which is preferable for logistical reasons) or two (which is preferable for academic reasons).

Conclusion

While the course has certainly been a success, it is hard to quantify the degree to which its aims have been met. Early indications are that students' progress in our graduate program has, on

---

[1]To preserve some degree of privacy, the presentation to which each set of comments pertained was not identified, and the order was randomized.

average, accelerated somewhat—though other factors may have contributed to this too. Students are generally better informed about what is going on in the Department than before. An unstated goal was to get students into the habit of voluntarily attending seminars regularly, and this has not been completely successful. Finally, our students' research skills do seem to have been improved by the course, though this is an impression that is hard to justify objectively.

The course has been interesting to develop and fun to teach. It has run stably for the past three years, and the material in it has slowly evolved as I discovered more useful information on skills for the professional researcher. One possible improvement is to end the course with a one-day conference organized by the students, with refereeing and presentation of thesis proposals. I believe that this kind of course provides a valuable means of strengthening a graduate program in computer science.

References

Anonymous. The publication game: beyond quality in the search for a lengthy vitae. *J Social Behavior and Personality 2*, 1 (1987), 3–12; and comments on the paper on pp.13–22 of the same issue.

Barzun, J. and Graff, H.F. *The modern researcher*. Harcourt Brace Jovanovich, New York, 1977 (third edition).

Brooks, F.P. *The mythical man-month*. Addison-Wesley, Reading, MA, 1978.

Bundy, A., du Boulay, B., Howe, J. and Plotkin, G. The researchers' bible. Teaching Paper No.4, Department of Artificial Intelligence, University of Edinburgh, Scotland, 1986.

Campbell, W.G. and Ballou, S.V. *Form and style: theses, reports, term papers*. Houghton Mifflin, Boston, MA, 1974.

Cleese, J. *Please can we have that the right way round?* Produced by Video Arts (20 min 16 mm film); distributed by International Telefilm Enterprises, 1976.

Crawford, S. and Stucki, L. Peer review and the changing research record. *J American Society for Information Science 41*, 3 (1990), 223–228.

Crutcher, K.A. How to succeed in science. *Perspectives in Biology and Medicine 23*, 2 (Winter 1991), 213–218.

Knuth, D.E., Larrabee, T. and Roberts, P.M. Mathematical writing. Report STAN-CS-88-1193, Department of Computer Science, Stanford University, Stanford, CA, 1988.

Kuhn, T.S. *The structure of scientific revolutions.* University of Chicago Press, Chicago, IL, 1970.

Lamport, L. *LATEX: user's guide and reference manual.* Addison-Wesley, Reading, MA, 1986.

Langley, P. Editorial: Advice to *Machine Learning* authors. *Machine Learning 5* (1990), 233–237.

Mental, X.Perry How to make a scientific lecture unbearable. *J Irreproducible Results 18*, 3 (March 1971).

Michaelson, H.B. How to present a paper orally. Chapter 24 of *How to write and publish engineering papers and reports.* ISI Press, Philadelphia, 1982.

Newton-Smith, W.H. *The rationality of science.* Routledge and Kegan Paul, Boston, MA, 1981.

Parberry, I. How to present a paper in theoretical computer science: a speaker's guide for students. *Sigact News 19*, 2 (Summer 1988), 42–47.

Parberry, I. A guide for new referees in theoretical computer science. *Sigact News 20*, 4 (Fall 1989), 92–109.

Sides, C.H. *How to write papers and reports about computer technology.* ISI Press, Philadelphia, 1984.

Tufte, E.R. *The visual display of quantitative information.* Graphics Press, Cheshire, CT, 1983.

Tukey, J.W. *Exploratory data analysis.* Addison-Wesley, Reading, Massachusetts, 1977.

University of Chicago Press *The Chicago manual of style (13th edition).* University of Chicago Press, Chicago, IL, 1982.

van Leunen, M-C. *A handbook for scholars.* Alfred A. Knopf, New York, 1986.

Williams, R. *The Mac is not a typewriter.* Peachpit Press, Berkeley, CA, 1990.

Witten, I.H. How to get a research grant. *Canadian Artificial Intelligence 24* (July 1990), 9–13.

Week	Responsibility	Activity	Assignment due
1	Instructor	Introduction to the course	
	Head of Department	Departmental expectations of graduate students	
2	Faculty member	Research seminar	
	Instructor	Discussion of Bundy *et al.* (1986)	
3	Technical staff	Departmental computing environment	
	Instructor	Document preparation tools	
4	Faculty member	Research seminar	
	Instructor	Computer Science literature, and the library	
5	Faculty member	Research seminar	
	Movie	Giving presentations	
6	Faculty member	Research seminar	
	Instructor	Discussion of seminar presentation styles	
7	Faculty member	Research seminar	
	Instructor	Writing critical reviews	
8	Faculty member	Research seminar	
	Student	Brief presentation 1	
9	Senior grad student	Research seminar	
	Student	Brief presentation 2	
10	Student	Brief presentation 3	
	Student	Brief presentation 4	
	Student	Brief presentation 5	
11	Student	Brief presentation 6	Review 1
	Student	Brief presentation 7	
	Student	Brief presentation 8	
12	Student	Brief presentation 9	Seminar digests
	Student	Brief presentation 10	
	Student	Brief presentation 11	
13	Faculty member	Research seminar	
	Invited guest	Brainstorming techniques	
14	Instructor	Preparing research proposals	Review 2
	Invited guest	Visual display of quantitative information	
15	Student	Full-length presentation 1	Thesis proposal
	Student	Full-length presentation 2	synopsis
16	Student	Full-length presentation 3	
	Student	Full-length presentation 4	
17	Senior grad student	Research seminar	
	Instructor	Writing academic resumes	
18	Instructor	Discussion of students' ideas for thesis proposals	
	Instructor	Refereeing	
19	Student	Full-length presentation 5	
	Student	Full-length presentation 6	
20	Student	Full-length presentation 7	
	Invited guest	Symbolic algebra package	
21	Student	Full-length presentation 8	
	Student	Full-length presentation 9	
22	Student	Full-length presentation 10	
	Student	Full-length presentation 11	Thesis proposal
23	Instructor	Thesis orals	
	Instructor	The publication business	
24	Instructor	Mock thesis oral	Examiner's report
25	Instructor	Science and ethics	
	Instructor	General discussion of course	

Table 1 Typical schedule for the course

Brief presentation	10%
Full-length presentation (including handout)	20%
Reviews (2)	25%
Thesis proposal synopsis	5%
Thesis proposal	25%
Examiners' report and mock oral	5%
Seminar digests and participation	10%

Table 2 Assessment