

An inquiry-based learning approach to teaching information retrieval

Gareth J. F. Jones

Received: 12 September 2007 / Accepted: 2 January 2009 / Published online: 22 January 2009
© Springer Science+Business Media, LLC 2009

Abstract The study of information retrieval (IR) has increased in interest and importance with the explosive growth of online information in recent years. Learning about IR within formal courses of study enables users of search engines to use them more knowledgeably and effectively, while providing the starting point for the explorations of new researchers into novel search technologies. Although IR can be taught in a traditional manner of formal classroom instruction with students being led through the details of the subject and expected to reproduce this in assessment, the nature of IR as a topic makes it an ideal subject for inquiry-based learning approaches to teaching. In an inquiry-based learning approach students are introduced to the principles of a subject and then encouraged to develop their understanding by solving structured or open problems. Working through solutions in subsequent class discussions enables students to appreciate the availability of alternative solutions as proposed by their classmates. Following this approach students not only learn the details of IR techniques, but significantly, naturally learn to apply them in solution of problems. In doing this they not only gain an appreciation of alternative solutions to a problem, but also how to assess their relative strengths and weaknesses. Developing confidence and skills in problem solving enables student assessment to be structured around solution of problems. Thus students can be assessed on the basis of their understanding and ability to apply techniques, rather simply their skill at reciting facts. This has the additional benefit of encouraging general problem solving skills which can be of benefit in other subjects. This approach to teaching IR was successfully implemented in an undergraduate module where students were assessed in a written examination exploring their knowledge and understanding of the principles of IR and their ability to apply them to solving problems, and a written assignment based on developing an individual research proposal.

Keywords Inquiry-based learning · Teaching information retrieval · Language technology integration

G. J. F. Jones (✉)
School of Computing, Dublin City University, Dublin 9, Ireland
e-mail: gjones@computing.dcu.ie; Gareth.Jones@computing.dcu.ie

1 Introduction

Interest in information retrieval (IR) as a subject of study has increased significantly in recent years. This has been driven by the very rapid growth in largely unstructured online information repositories, principally World Wide Web pages, but increasingly other digital media sources such as audio and video, and the need to be able to efficiently locate information relevant to a user's information need from within these collections. Until recently IR as a subject featured primarily in courses for librarians and information scientists, but is now appearing at undergraduate and graduate level in many courses in computing and information technology, and indeed in other courses where the ability to use search engines effectively is seen as important. The students taking these courses have differing backgrounds and learning requirements, and IR modules must be tailored appropriately to select between the wide range of topics that it encompasses.

Beyond formal instruction, the new challenges and opportunities for IR technologies arising from varying user needs, expectations and expertise, and the available information sources, mean that IR has emerged as a dynamic and rapidly developing subject with a vibrant and growing research community. However, despite the proliferation of new areas of IR research such as bioinformatics, question answering, multimedia IR, topic tracking and web search, the fundamental issue of IR remains essentially the same, namely satisfying the information need of a user expressed through some form of search request.

For some of those studying IR, it is sufficient to learn how to use commercial search engines effectively (Lazarinis 2007). For others it is important for them to move beyond this to a technical understanding of IR. These students, in order to fully appreciate, exploit and contribute to advances in search technologies students, need to acquire a sound knowledge of an appropriate set of fundamental issues and techniques, open questions and relevant research strategies. The topics included in a complete introduction to IR should include text processing, but might also introduce indexing and search for multimedia content, content summarization, as applied in snippet generation, and information extraction, as used in question answering. Learning about IR methods and related technologies enables users of search engines to use them more knowledgeably and effectively, while providing the starting point for the studies of new IR researchers. However, the nature of IR as a topic also makes it an ideal subject for developing a range of interdisciplinary and transferrable skills in those studying it. In many ways IR is no different to other subjects, however the author would argue that the basic principles of IR are sufficiently easy to grasp, and in combination with the various practical challenges of information searching, it forms an ideal subject for exploring creative teaching, learning and assessment methods.

This paper reviews some of the author's experiences in developing, delivering and assessing an undergraduate module in IR technologies. It relates these to studies in student learning, expectations and assessment, gives some observations on student prerequisites, and demonstrates the potential for original research contributions in student projects. While the author's experiences are restricted to student groups studying computationally focused hard science degrees, the approach taken here is generally qualitative and relies only on a basic level of mathematics skills, meaning that the approach explored in this paper could, he believes, generally be applied with similar results to students taking softer science courses in information or library sciences.

The remainder of this paper is organised as follows: Sect. 2 considers the content of an undergraduate module in IR, Sect. 3 reviews some relevant studies in student learning, Sect. 4 describes the design and outcomes of a final year Bachelors module in Information

Access developed by the author some years ago, and Sect. 5 concludes with some remarks on the outputs from this module and how it might be extended.

2 What should be covered by a module in information retrieval?

Within computing degrees IR is typically taken as a one semester module within final year undergraduate and masters programmes. Students may of course have studied related topics in other modules, such as multimedia technologies, the internet or computer networks in general, relevant topics in mathematics, or one or more topics in natural language processing. However, such related modules are often optional and knowledge of all of these relevant topics cannot be assumed of the whole cohort opting to take an IR module. Within a programme of study where these topics may have been covered elsewhere by some of the class, the IR module must necessarily repeat some material, but must seek to present and address it in an alternative way relevant to IR. The issue of the background of students taking modules in IR is generally problematic. This relates not just to pre-existing knowledge of the students, but also as MacFarlane (2007) notes with regard to the need to teach the necessary elements of mathematics, also to student motivation and confidence in individual disciplines. The large number and breadth of topics which can be defined within the study of IR mean that instructors must take account of the expected background and required learning outcomes of the cohort to be taught, and design a course including the necessary topics at an appropriate depth. This represents a difficult challenge, and it is easy to overestimate or underestimate the capabilities of a particular class.

More fundamentally, since IR is only a single semester module, it is in the author's view unrealistic to seek to teach both the basic concepts and give detailed coverage of a wide range of the related techniques and algorithms. In the author's experience, it is all too easy for instructors to adopt, probably without considering the implications, a strategy of determining that students must be exposed to every topic related to a subject in as much detail as possible in the time available, since all this material is "vital" to an acceptable working knowledge of the subject. Of course, the danger of doing this is that the module becomes so packed with content that in order to cover everything, the instructor finds it necessary to deliver it as fast as realistically possible in classical lecture delivery mode where students are expected to sit and absorb as many facts about contemporary IR methods as possible. In choosing how to structure an IR module, the instructor could ask themselves, should the students be exposed to a full detailed treatment of data structures for efficient search or all the variations on "tf*idf" weighting, or is it sufficient to cover the basic principles with some simple and effective solutions and let students explore the application of these techniques in practical problems? If one steps back to consider this scenario for a moment, it is clear that it is easy to lose sight of both the objectives of effective teaching and learning from the students' perspective (Bligh 1998), and also risks the students not being able to see the wood for the trees in terms of appreciating the fundamental issues of IR amidst the details of current techniques which may quickly be replaced within a year or two. The topic of curriculum design for different student groupings and the limits imposed by the teaching time available is explored in more detail in (Bawden et al. 2007).

Taking these issues together, some years ago while working at the University of Exeter in the U.K., the author developed a module which he called *Information Access*. While the expression Information Access may have differing interpretations for some readers, for example referring to copyright, privacy or security; it is used here to describe a more

general concept than that which is usually associated with IR. The Information Access module encompassed an introduction to IR combined with a range of technologies closely related to IR which can support information seeking functionality for a range of emerging systems. The idea was to look at technologies enabling access to information in a more general sense than the process of retrieving it. Before describing the design and methodology of the Information Access module, the next section briefly outlines some of its underlying motivations based on research in student learning. Much of this review is derived from Bligh (1998) which provides an excellent summary of student learning taken from a wide range of sources.

3 Teaching and learning methods and information retrieval

Students can study and learn any subject at different levels, and this is no less true of IR. This section briefly summarizes some of the key issues in student learning and then relates these to the teaching and assessment of IR.

3.1 Student learning modes and assessment

Student learning can be generally classified into two forms: *surface* (or shallow) and *deep* (Fry et al. 1999). Students engaging in surface learning are generally found to be attempting to store information without analyzing and reflecting on it. In this mode of study, learning is seen as a process of acquiring facts related to a subject, and learning the principles and procedures associated with them. When following this style of learning, students in general see the role of the instructor as being to present information and the student's role as to reproduce this information in their assessments in order to demonstrate that they know it. Instruction in this manner is often referred to as the classical lecture presentation process. This involves large amounts of rote learning on the part of the student, and fairly unchallenging lectures from the perspective of both the lecturer, who presents a prepared lecture "script", and the student who is expected to listen to (and it is generally assumed "understand") the material without actually doing anything else. By contrast, in deep learning the emphasis is much more on thought and reflection rather than memorization of facts. Rather than learning by rote the "words of wisdom" from their expert lecturer, students actively integrate new ideas with those already possessed. Students using a deep approach look for the fundamental principles associated with the subject. They distinguish the principles of the subject from examples which demonstrate these principles in applications and are then able to exercise imagination within the subject.

Traditional lectures often encourage the surface approach with students listening to the instructor delivering facts about the subject at hand, and taking notes or working through handouts which they then learn largely verbatim for examinations. This mode of surface study actually discourages the very intellectual skills of thinking, integration and imagination that higher education claims to foster in students. While one may doubt that this rather poor form of study is a widespread reality, the author has encountered at first hand examples of questions in degree examinations which are merely the recitation of the derivation of mathematic equations covered in the course. Such questions convey nothing to the examiner in terms of the candidate's understanding of the material or their ability to apply it in solution of a real-world problem. Even if the questions require analysis and application of material to solve them, students may still attempt to "solve" them by reproducing lecture material by rote. A number of the author's colleagues have commented

over the years on students who, when unable to properly interpret or answer a question, attempt to pick up marks by reciting sections of their notes which they think may be relevant to the question in the hope that the examiner will find something relevant in their answer. Studying in this manner while less intellectually engaging can actually often involve more study time and effort in the long term for the student as they attempt to memorize everything, rather than spending time understanding the principles in the first place, and then seeking to apply them to problems set for assessment. It is often assumed, or perhaps more accurately hoped, by instructors, that students will make the time to reflect on material delivered in lectures and explore it afterwards in private study. But high teaching and course assessment loads associated with module assignments and examinations, which are often driven by factual regurgitation, mean that this often does not happen.

Students and indeed lecturers often find classical lectures rather boring. Lecturers are simply telling students what they already know about the subject, and students receive the information in an entirely passive mode. In this situation students may be tempted to miss classes where they can copy notes from colleagues or download them from a module website. Students perceive no apparent added value from attendance at the lectures. The lecture experience can become a more engaging and valuable experience for all parties when they become interactive. In the author's own classes he has observed that interest increases and learning is often improved if students are asked questions in class rather than told facts. However, delivering lectures in this form requires more effort on the part of the lecturer, who must be prepared to do more than recite information and be prepared to challenge and engage with students, and also on the part of the students, who must really listen, engage and think during the class, rather than sit and copy down notes or frequently perhaps start daydreaming.

The objective of challenging students in this way in class is to encourage them to apply basic concepts to problems and move towards deep learning. Working in this way students can be led towards developing the key concepts for themselves, as well as being able to learn to analyze positive and negative features of the methods covered and thinking critically about proposed solutions to problems. This approach to teaching is referred to as *inquiry-based*, or sometimes, not quite correctly, *problem-based* learning. This is sometimes viewed as students developing all details of a subject for themselves through interactive inquiry or by solving a problem without instruction. It has been argued that since this is often quite unrealistic, that this mode of learning should be avoided (Kirschner et al. 2006). However, this argument confuses inquiry-based learning with unguided *discovery-based* learning, and the picture is rather different if one uses inquiry-based methods within an integrated framework of study combining direct instruction to form a scaffold of the subject with working with problems to develop deeper understanding and soft-skills such as self-directed study (Hmelo-Silver et al. 2007). Thus the fundamental principles of the subject need to be established to enable students to build on these in subsequent carefully guided self-discovery. The balance or trade-off between formal instruction and inquiry-based teaching is one of the specific challenges of this method of instruction. If inquiry-based learning is to work successfully the instructor needs to be aware of this, and to actively adapt the structure and form of a teaching session as it proceeds.

It is frequently said that students are motivated by assessment (McFarlane 2005). Thus if they can see the relevance of the material being covered to the module assessment, they are more likely to be motivated to engage with it in the appropriate manner. The desire to achieve generally improves motivation and learning, but students need to know what is to be achieved! Thus it is important to tell them the objectives of the module and each lecture at the start, so that they know what they should learn from it. In terms of assessment, one

can rely on reciting information to find out what facts the students know, but if they have actively engaged with the material in classes and have been encouraged to approach the material from an imaginative and creative perspective, this can be pursued in the assessment as well. This leads to the opportunity for the examiner to ask questions which explore the candidates understanding of the principles of the subject and apply it to novel problems, rather than reproducing lecture notes or filling in the gaps in small variations in examples taken from lecture material. For example, students may be presented with a novel practical scenario which they need to analyze and then select and integrate material from the course to solve them.

3.2 Teaching and learning information retrieval

Information retrieval is a subject built around fundamental principles which are generally accessible to students of a wide range of skills and abilities. In addition, the rapidly developing range of technologies associated with IR mean also that it is a subject in which imagination plays a key role in new developments. IR modules thus represent a great opportunity to be structured around inquiry-based learning since it can enable students to demonstrate their imagination and creativity, and also encourage them to develop a deep approach to learning. It can also be used as a means for them to acquire subject independent skills for learning, and through careful design of assessment, to practise and demonstrate their command of these skills. For example, to extract, summarize and contrast material from their reading and report this in clear, properly documented technical prose. In order to encourage students to adopt a deep learning approach, the lecturer must “lecture“ less, convince students of the intellectual aims of their course, and create opportunities, in classes and outside, in which thinking and reflection can flourish.

Another positive aspect of IR as a topic of study is that if students can see the relevance of a module to their lives it generally aids their motivation (Merrill 2007). Most students taking IR courses are regular users of search engines and digital technologies in general, making the relevance of the material clear to their lives is thus not generally a problem for the instructor. Mizzaro (2007) notes that students can be more engaged in their study of IR if its relevance to their own experience of web search is made clear. This can be achieved for example by setting them search tasks outside class, and asking them to report back on their findings. Further examples can be found in areas such as search of their own digital photo collections, where the failure of current applications to support search means that large numbers of photos are never looked at again once they have been archived to their computer hard drive.

In terms of assessment, it is relatively easy to set traditional questions requiring students merely to memorize content. They either learn the script delivered in the lectures, where they take down their own notes verbatim, or use handouts made available either in printed form, or more likely, made available for download from a module webpage. However, as we have seen, this approach fails many of the generally desired learning outcomes of a university level education. Rather than merely measuring the extent to which a candidate has memorized material, this requires examiners to develop assessments which, while establishing that the fundamentals of a subject are known, examine the student’s understanding. One way to do this is to set problem-based questions related to, but not identical to, those covered in classes.

Considering assessment in this way is consistent with the approach of *outcome-based learning* (Furman 1994). In this strategy for assessment, measurable objectives of student learning are established for the learning unit as part of its development and the assessment

is then designed to measure the degree to which these have been achieved. While these may of course be implicit in the design and assessment of modules designed by other means, making them explicit in the module description and the planning of the assessment makes it clear to all parties what is to be assessed.

In delivering an IR module structured around inquiry-based learning with stated learning outcomes, the author experimented with providing copies of lecture slides to students in advance of lectures, enabling students to bring the slides with them to the class so that they could annotate them with personal additional comments, or making the slides available after the lecture, to encourage students to listen and interact with the material during the lecture. Students, perhaps unsurprisingly, in the author's experience universally favoured being provided with the materials in advance. While providing notes in this way seems to be pedagogically justified if the students make the intended use of them; it inevitably also creates problems, students who attend the class with the notes may not pay full attention, since they may feel that they already have the content of the lecture in these provided notes, and students who are inclined to skip lectures can do so in the knowledge that the "lecture" is "on the web" (or so they believe!). Depending on the content of the notes, and the correlation between them and the content of the lecture itself, the added value of actually attending or even participating in the lectures, and the assessment methods used in examinations or coursework, the students may feel fully justified in missing the class. There is after all little point in spending time in a class if the student can just read it up from notes provided for the lecture in advance of the assessment and still be able to gain a distinction level mark. One could argue slightly dogmatically that students must attend classes, but this is academically and intellectually difficult to defend if there is visible evidence of the assessment outcomes showing there to be no benefit arising from attendance. The challenge then to the instructor is to make lectures relevant, interesting and engaging, and to seek to ensure that attendance at classes adds value that cannot be gained by reading up on the subject in private study.

Furthermore, if continuous assessment assignments are based around reviews of existing work, e.g. writing review essays, many students take the approach of writing submissions that are highly derivative of recommended reading. Even if writing essays of this sort does not constitute plagiarism, and in some cases it probably does, it is not at all clear that students really gain much from such assignments. The value of these assessments is potentially much greater if they seek to establish that the fundamentals of the subject have been understood, and then to explore the student's ability to break down a problem and make use of their knowledge of the principles and techniques of the subject to address and report it in a creative way.

In summary an assessment should seek to find out to what degree the key principles of the subject are known to the student, and assess their ability to understand these in the context of a problem scenario and to select and combine appropriate technologies to provide well planned and justified solutions to this problem. Suitable assessments can be made using both written examinations and coursework assignments. In the author's experience, as well as examining the student's problem solving and reporting skills, the use of problem-based assessments greatly reduces the scope for simple reproduction of existing work and makes the student's approach to analyzing and solving the problem visible to the examiner. While the utility of inquiry-based learning relative to more standard teaching methods is still an active debate (Kirschner et al. 2006; Hmelo-Silver et al. 2007), the author's experience of using it has been very positive in terms of student engagement with material during teaching and assessment, and the demonstrated learning outcomes.

4 Information access

4.1 Background

While working at the University of Exeter, the author developed an undergraduate module which he called *Information Access*. Exeter was active in promoting inquiry-based study and outcome-based learning, and teaching staff were required to incorporate these methods in module design and subsequent assessments. The Information Access module was aimed primarily at final year bachelor students studying computer science or cognitive science, but was also taken by small numbers of other students, predominantly visiting students from universities in mainland Europe. The class had between 50 and 60 students each year. The module took IR as its hub, but introduced a range of related technologies for information indexing, searching and presentation within unstructured document collections, and explored their integration to address information access tasks. The module assumed a prerequisite of basic undergraduate statistics and introductory artificial intelligence methods. In line with the university's education strategy, the general philosophy of teaching and learning for these students within their programmes of study was very much based on inquiry-based analysis of problem scenarios, and design and implementation of solutions. The students were thus used to being challenged to learn a new subject by solving problems.

The technical focus of the module was to establish the unchanging issues and challenges of information access tasks centred around IR related applications. In practice this was principally to convey the concepts of user information need, document collections, uncertainty of relevance and to make clear why accurate and reliable IR is difficult! The module introduced current techniques from IR very much from a practical rather than a highly theoretical perspective, to enable students to build effective prototype tools for IR. For example, they developed a good conceptual understanding of well established issues of document ranking algorithms, and how these are applied in probabilistic IR without needing to study the underlying theory of the probabilistic model. Students completing the module were expected to be able to follow a "recipe" for the construction of an effective small scale IR system, for example to follow easily the description of the BM25 model described in (Robertson and Spärck Jones 2006).

Another aspect in design and delivery of the module was to ensure the students understood that IR is a rapidly developing subject within which there is scope for them to make novel and significant personal contributions. While this last aspect is perhaps obvious to experienced researchers, this is very much not the case for many undergraduate students. Students often do not realize that computing is a live subject within which they are free to propose, test and report new ideas; and that problems often do not have clear cut right and wrong answers, but have alternative solutions with contrasting strengths and weaknesses which must be taken into account when selecting from among them. As part of the emphasis on the development of novel technologies, the module also introduced the need for effective evaluation of IR applications to enable testing of ideas and, where appropriate, comparison of alternative potential solutions.

4.2 Syllabus

The intention of the module was to give students a broad general understanding of the issues of indexing, identification of relevant information and its presentation, and how this can be achieved by the integration of a diverse range of technologies. To this end, the

module began by introducing standard IR topics including stop word removal, stemming, file structures, Boolean and ranked retrieval and term weighting. It then introduced the related information access topics of hypertext, information extraction, machine translation, speech recognition, information visualization, intelligent agents and summarization. Coverage of each topic introduced relevant evaluation metrics and methods for using them. For example in the case of IR, this introduced the Cranfield laboratory evaluation paradigm as used at TREC, giving the definitions of *precision* and *recall*, and the design of a suitable test collection. This enabled classes to explore a wide range of integrated information access scenarios, including topics such as cross-language IR, spoken document retrieval (SDR), information exploration using graphical visualization, question answering, agent-based information discovery and delivery, and web searching. It should be emphasized that each topic was covered at an introductory level of definition, establishing the fundamental challenges and problems, and in outline current methods used to implement each one. For example, in the case of speech recognition for SDR, the module reviewed the state-of-the-art in speech recognition technologies, established that they produce errorful output, asked why this is so, explained why this will continue to be the case for the foreseeable future, and then considered whether the noisy output from a speech recognition system can be used for effective SDR. This discussion was then extended to consider whether we can measure the impact of transcription errors on retrieval effectiveness and the how specific types of transcription error will affect the parameters and behaviour of an IR system.

4.3 Teaching materials and delivery

The wide range of topics covered meant that there was no suitable single set text that could be used. However, the module design sought to turn this into an opportunity for the students to develop information searching skills of their own. Key texts on each of the module topics were identified and made available in the university library. In addition key research papers were identified, some of these were tutorial style papers on the topics, while others represented examples of current research combining topics covered in the module. Copies of these papers were made available to the students as a module “reader” in the School library. Students were able to make copies of these materials for private study.

Students were provided with several online handouts at the start of the module. These comprised an introductory overview of the objectives and teaching approach of the module, and a detailed list of the papers and texts provided for the module. The assessment elements and methods were explained, along with the expectations for reading of the module materials. For each topic the lectures introduced the principles of the topic, but as often as possible students were encouraged to think critically about the issues relating to a topic and to try to recognize the inherent challenges in each topic for themselves. The notes used in each lecture were also made available online. As discussed earlier there seems to be no ideal approach to the provision and timing of the distribution of lecture notes. After some experimentation it was generally found to be most effective to make handouts available in advance of classes, but to keep them fairly brief and make clear that they represented a minimum requirement of knowledge on the topic, and that students were expected to read around the topics from the provided books and research papers.

One or two of the provided papers were set as readings for each week of the module delivery, and it was made clear that students would be expected to be familiar with these materials in lectures in following weeks. The reasoning here was that students at this stage of their study are generally not familiar with reading research materials, by setting specific

readings they could start gaining familiarity with this style of writing. The more extensive list of publications available in the module reader gave them a starting point for wider exploration, and also to help them identify leading researchers and venues to look for further IR research publications. They were further encouraged to read beyond this list both by following references and searching the web as part of the module continuous assessment assignment.

The general structure of delivery of the Information Access module was largely consistent with that recently described by Merrill (2007) which advocates a task-centered instructional strategy. This is based around his “Pebble-in-the-Pond” model introduced in earlier work (Merrill 2002). This strategy advocates a progression of introducing a real-world task relating to the subject, and then a progression to others illustrating related or more complex task situations. These are then broken down into their components which are described in the context to which they will be applied.

The module was motivated to the class in the first lecture by introducing the topic of search from unstructured information sources and the associated challenges. Consistent with Merrill’s proposals, this was contextualized using examples such as retrieval from a very large item set such as the web and the multiple challenges related to multilingual search across multimedia data. Following this the topic of IR was covered in around four lectures. This introduced the key concepts such as *information need* and *relevance*, the technologies such as *stemming* and *term weighting*, and then evaluation. Each of the other topics in the module, e.g. machine translation, summarization, etc. were introduced more briefly in two to three lectures. The first of these introduced the topic and its challenges, and the second covered a relevant recently published research study which combined multiple technical topics. Within these sessions as well as the direct instruction discussed previously, creative thinking was encouraged by posing questions in a progressive manner as the description of the topic proceeded. For example, what are the options when choosing the contents of a summary, and, once these had been identified, what are the advantages and disadvantages of different approaches to doing this? Teaching sessions also looked at experiments reported in the published research studied in the classes to identify the hypotheses under investigation, to question the validity of the claims made and consider whether the experiments might usefully be extended or the results analyzed in different ways. Once these had been established more open questions were posed. For example, in order to achieve a particular IR task, such as cross language IR, the class looked at questions such as, what are the issues beyond monolingual IR that must be addressed? how might these be addressed in a practical system? what are the strengths, limitations, etc. of particular potential solutions? and how might the effectiveness of the proposed methods be tested? Interaction with the class was based on several approaches including simple direct questions to the class as a whole, more open questions to the class leading to discussion, and more extensive questions in the form of partial or complete problem scenarios for which the class were asked to work in pairs or small groups to propose answers which were then reported to and discussed by the whole class.

4.4 Assessment

The module was assessed 80% by examination and 20% by a written continuous assessment assignment. The ratio of the split between examination and continuous assessment was typical of the modules taken by these students. Variations from this norm had to be approved by the course teaching committee. While a special case could have been made to this committee to change this split in favour of a greater weighting of the continuous

assessment, overall it was found that the results were in general a fair reflection of candidates subject knowledge and relevant problem solving skills.

The examination was of a fairly traditional structure with free choice of three questions from five in 2 hours. Individual questions combined multiple topics from the module, reflecting the inherently integrated nature of the information access problems addressed in the module. The method of teaching and the general preparation of this class meant that it was possible to structure many of the examination questions around problems and scenarios. The beginning of each question required students to explain basic definitions or identify key issues. Questions then moved to problem-based questioning where candidates needed to analyze a given scenario and to develop creative solutions by selecting and applying appropriate technologies introduced in lectures or that they had found in independent reading or they were free to propose their own ideas as solutions to the problem, and then to suggest how their solutions might be evaluated based on experience gained during the module.

Although not universally popular among the students, the combination of topics within individual questions meant that candidates were encouraged to study the whole range of topics presented in the module since they were very unlikely to perform well if they only engaged in selective revision of parts of the syllabus.

The more novel element of the module assessment was the continuous assessment assignment. The objective here was to explore the students' ability to examine a problem, and to propose a technically realistic solution and its evaluation, and their ability to report this. Students were asked to select and consider one of a number of given "research" problems in information access. The problems were designed to generally be quite novel information access scenarios for which there not ready made solutions in the current research literature. They were required to report their solution as a formally structured research paper. A template of a standard research paper format was explained: abstract, introduction, literature review, proposal, method of assessing proposal, and anticipated possible results and conclusion, with properly formatted references. Submissions were required to be in this format with a prescribed maximum word limit of 3,000 words. This form of assessment exercised a number of important transferrable skills, as well as testing understanding of IR and related technologies. The abstract tested students' ability to construct and write a succinct summary of a document. The introduction needed them to give suitable background and motivation, and detail of the topic covered and paper structure. The review required students to select relevant material and ignore non-relevant material from reviewed documents contrasting related work as appropriate, and then concisely express this material in an integrated fashion leading to the justification of a research proposal. A means of implementing and then evaluating the proposal had to be described, and students needed to explore anticipated results and conclusions. This required a clear and precise practical explanation and justification for the technologies selected and how they could be integrated to solve the problem, and the inclusion of an often novel evaluation framework with suggestions for suitable data sources and evaluation metrics. The reference section had to properly cite reviewed research papers, credit was given for wide reading of materials beyond that introduced in the lectures.

A marking scheme making clear the requirements necessary of an ideal submission, and progressively weaker ones associated with each grade was included with the assignment. A sample of the assignment marking scheme is given in the [Appendix](#). It was the policy of the department to provide written personal feedback on student assignments. Although time consuming and expensive to operate this policy of providing feedback in this way was highly regarded by students. The author thus provided students with individual critical and

constructive feedback on their work. Using this published marking scheme and written feedback to the students, the author received only one query about the assigned mark out of more than 200 students who submitted assignments during the years that the module was delivered.

Students often found this assignment a very challenging exercise for a variety of reasons. Most had never attempted to write a paper of this type before. Differing aspects challenged individual students, for example writing a meaningful abstract, providing a sufficiently detailed review within a tight word limit, developing new ideas, or considering how to evaluate their ideas. Ultimately in informal feedback students generally agreed that writing the assignment following the research paper template formed a very useful learning exercise. From an assessment perspective one particular strength of this assignment over a standard essay type review, is that students could not simply restate standard materials from books and papers. The submissions were generally more interesting to read when marking them than standard essays, and it was generally clear where material was taken from existing sources since the technical knowledge displayed and the writing style were significantly different from the student's own work in other parts of the paper.

4.5 Exploiting module outputs in student projects

Each year a number of students completing the Information Access module went on to undertake final year projects under the author's supervision. The best of these were accepted for publication at international conferences (Jones et al. 2000; Jones and Gabb 2002; Jones and Edens 2002), a number of other students completed excellent projects which were not submitted for publication. The author found that the module provided excellent preparation for these projects in terms of basic subject knowledge, but also the student's ability to conduct background research, creative thinking, and evaluation.

5 Concluding remarks

The module in Information Access was delivered at the University of Exeter for 5 years prior to the author's move to Dublin City University. It proved consistently popular as an option choice and feedback from students indicated that many of them found IR engaging and interesting. Overall the author believes that it achieved its learning objectives.

What other components might be incorporated in a module of this type to improve student learning with respect to both IR and skills more generally? One request from the students was for practical exercises where they would be able to explore the behaviour of the technologies covered in the module. Unfortunately, it was not possible to resource the development of suitable materials and staffing of practical classes on this scale. Incorporating practical work to improve understanding of ideas covered in the module would almost certainly have enhanced the student learning experience. The continuous assessment exercise might even be extended to ask students to implement and test their ideas. Although if the module were to be extended to include these practical elements, its overall weighting in the students' course would need to be increased.

Another possibility could be to introduce student presentations, either to describe course materials, reviews of recent research or perhaps the outcomes of the continuous assessment exercise. Mizzaro (2007) notes positive reactions from his students to the use of student presentations in his IR course. One of the author's former colleagues at Exeter received similarly positive responses to the extensive use of student presentations in a module on

artificial intelligence available as another option to the student groups taking the Information Access module.

Teaching IR or any other subject using an inquiry-based approach requires that students know how to respond to this environment. While strong students will generally respond well to any approach to teaching, other students can find it confusing if the mode of teaching and assessment changes radically from that with which they are familiar, since it represents a very significant challenge to the study and learning skills they have developed previously. For example, in rote learning and reproduction of this material in assessment. Attempting to develop inquiry-based learning skills in students not previously exposed to them while teaching IR within a single module is perhaps too much of a challenge for both students and their instructor. In the author's experience it serves students well to introduce inquiry-based teaching early in their studies, for example within an introductory module in information systems in their first year of study.

Appendix

Marking criteria

The following can be said of a perfect submission:

“You have written an excellent and well structured paper. The writing in the paper is clear and concise. You have provided full reference details of your sources and you show clear evidence of reading beyond material provided in lectures. You have shown excellent judgment in your choice of relevant material from your sources and smoothly integrated these into a clearly argued discussion of existing material. You have suggested well motivated methods by which existing techniques might be combined to address the research problem and commented on potential weaknesses and may have proposed some novel extension to these methods. You have designed a logically structured experimental research plan with clear justification of your design choices. You have demonstrated an excellent understanding of information system evaluation. An Information Access researcher would be able to use your CA to prepare and carry out the experimental plan described without needing to consult you or your background sources for any clarification.”

The following gives the criteria for assessment against this ideal:

Grade A: You have written a very good and well structured paper. You have documented your sources well and integrated them well into your discussions, including material beyond that covered in lectures. You have suggested some well considered means by which existing technologies might be integrated. You have included a complete experimental plan with justified design choices and described a very good evaluation strategy.

Grade B: Your paper is good but the structure may not be smooth or some points may not be well explained. You have shown good evidence of limited background reading. You have made a basic attempt to address the research problem. Your experimental design may not be complete or some decisions are not justified.

Grade C: Your paper is lacking in some areas or the flow of ideas is not logical. You have shown evidence of some background reading. Background material is included in a disjointed manner. Your research ideas are weak or not properly explained and justified. You have developed an experimental plan, but it has gaps or significant flaws in the design.

Grade D: You have written a complete paper, but it has significant faults in its structure. You have shown little sign of background reading and your discussion is very limited and difficult to follow. Your research ideas are very weak or may not be completely logical. There is some evidence of an experimental plan, but it would need significant revision before it could be carried out.

Grade U: Your paper is incomplete. There is no evidence of background reading. Your writing is poor and difficult to follow. Your experimental plan is incomplete and very poorly structured.

References

- Bawden, D., Bates, J., Steinerovu, J., Vakkari, P., & Vilar, P. (2007). Information retrieval curricula: Contexts and perspectives. In *Proceedings of the first international workshop on teaching and learning in information retrieval (TLIR 2007)* (pp. 55–60). London, UK. <http://www.bcs.org/server.php?show=ConWebDoc.8777>. Accessed 29 Dec 2008.
- Bligh, D. (1998). *What's the use of lectures?* (5th ed). Exeter: Intellect.
- Fry, H., Ketteridge, S., & Marshall, S. (1999). Understanding student learning. In H. Fry, S. Ketteridge, and S. Marshall (Eds.), *A handbook for teaching & learning in higher education* (pp. 21–40). London: Kogan Page.
- Furman, G. C. (1994). Outcome-based education and accountability. *Education and Urban Society*, 26(4), 417–437.
- Hmelo-Silver, C. E., Duncan, R. G. & Chinn, C. A. (2007). Scaffolding and achievement in problem-based and inquiry learning: A response to Kirschner, Sweller, Clark (2006). *Educational Psychologist*, 42(2), 99–107.
- Jones, G. J. F., & Edens, R. J. (2002). Automated alignment and annotation of audio-visual presentations. In *Proceedings of the sixth European conference on research and development for digital libraries (ECDL 2002)* (pp. 276–291). Rome, Italy.
- Jones, G. J. F., & Gabb, S. M. (2002). Visualisation tool for topic tracking analysis and development. In *Proceedings of the twenty-fifth annual international ACM SIGIR conference on research and development in information retrieval (SIGIR 2002)* (pp. 389–390). Tampere, Finland.
- Jones, G. J. F., Quedsted, D. J., & Thomson, K. E. (2000). Personalised delivery of news articles from multiple sources. In *Proceedings of the fourth European conference on digital libraries (ECDL 2000)* (pp. 340–343). Lisbon, Portugal.
- Kirschner, P. A., Sweller, J. & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, inquiry-based teaching. *Educational Psychologist*, 41(2), 75–86.
- Lazarinis, F. (2007). Forming an instructional approach to teach web searching skills to non-English users. *Program: Electronic Library and Information Systems*, 41(2), 170–179.
- MacFarlane, A. (2007). Pedagogic challenges in information retrieval—teaching mathematics to post-graduate information science students. In *Proceedings of the first international workshop on teaching and learning in information retrieval (TLIR 2007)* (pp. 39–44). London, UK. <http://www.bcs.org/server.php?show=ConWebDoc.8771>. Accessed 29 Dec 2008.
- McFarlane, K. (2005). Learning about employability. In P. Hartley, A. Woods, M. Pill (Eds.), *Enhancing teaching in higher education: New approaches for improving student learning* (Chapt. 12, pp. 133–144). New York: Routledge.
- Merrill, M. D. (2002). A pebble-in-the-pond model for instructional design. *Performance Improvement*, 41(7), 39–44.
- Merrill, M. D. (2007). A task-centered instructional strategy. *Journal of Research on Technology in Education*, 40(1), 33–50.
- Mizzaro, S. (2007). Teaching of web information retrieval: Web first or IR first?. In *Proceedings of the first international workshop on teaching and learning in information retrieval (TLIR 2007)* (pp. 50–54). London, UK. <http://www.bcs.org/server.php?show=ConWebDoc.8774>. Accessed 29 Dec 2008.
- Robertson, S., & Spärck Jones, K. (1994). Simple, proven approaches to text retrieval (Revised edition 1996, 1997, 2006). Technical Report 356. University of Cambridge Computer Laboratory.