

Simon P Bates*
 School of Physics
 University of Edinburgh
 Mayfield Road
 Edinburgh EH9 3JZ
 *s.p.bates@ed.ac.uk

Karen Howie
 and
 Alexander St J Murphy
 School of Biological
 Sciences
 University of Edinburgh
 Mayfield Road
 Edinburgh EH9 3JZ

In its simplest form, these questions may be employed to simply break-up the lecture, to regain audience focus and attention and as a mild diversion to the main business of the lecture time around halfway through.

The use of electronic voting systems in large group lectures: challenges and opportunities

Abstract

We describe pedagogical, technical and operational issues associated with the introduction of an electronic voting system into large first-year undergraduate lectures. The rationale for doing so is to transform the lecture experience from a one-way transmission of information in to a two-way conversation between lecturer and students, mediated by the technology. We discuss some of the logistics involved, such as choice of handset, cost and siting within a lecture theatre as well as the aspects of pedagogy, such as the requirements of a good question for these interactive episodes. We present a number of possible use scenarios and evaluate student and staff experiences of the process.

Introduction

"Despite the changes in the learning environment, teaching methods do not appear to have changed considerably. Initial findings from research suggest that many staff still see teaching primarily in terms of transmission of information, mainly through lectures."
 Dearing¹

The lecture is still the mainstay of university teaching. Its origins can be traced back in history, to when reading aloud to a group addressed the fundamental bottleneck in learning and teaching (the availability of books). Despite the enormous changes that have taken place over the last generation in terms of the size, diversity, expectations and career choices of the student cohort that come to University, the role of the lecture in scientific disciplines has remained largely unchanged. The traditional lecture, however inspirational it might be, is essentially a one-way transmission of information to the students, an exposition of the course material, mapped out along a carefully constructed A-Z of the course syllabus.

Concurrent with the shifting topology of the student landscape has been the explosion of computing and information technology to its present, almost ubiquitous state. The application of this within higher education has in general lagged behind social contexts. Compounding this lag is the fact that our students are now 'digital natives' – exposed to computing in education from an early age – whereas most of us (however enthusiastically we adopt) remain 'digital immigrants'². All of these factors add to the inertia that ensures, by and large, lectures continue to function in 'transmission' mode. This is especially true for large classes (>100 students). As Flowers has put it "Why Change? Been doin' it this way for 4000 years"³. Here, the pervasive influence of Information and Communication Technologies (ICT) can actually have a detrimental effect; a lecturer can easily make a presentation from a set of *PowerPoint* notes and after the lecture deploy the same notes on the World Wide Web. If the students perceive the net worth of the lecture as simply acquisition of the notes, and these notes are available in their entirety on the Web, they may well not attend.

The challenge is, therefore, to try to actively engage the students in the lecture, to develop it to be something more akin to a two-way conversation than a one-way transmission of information. Large classes present a particular problem here by virtue of their very size; one is simply precluded from striking up an interactive conversation with a hundred or more students.

This paper presents a review of one of the ways in which the use of technology can be used in a lecture context to mediate interactive engagement; via handheld, remote devices used to vote on questions, similar to systems popularised in TV shows such as *Who Wants To Be A Millionaire*. Based on our own recent experiences in Edinburgh, which over the last year has seen this electronic voting system used in three large first year undergraduate classes, we consider both practical and pedagogical issues

associated with incorporating this methodology into the curriculum. The paper is organised as follows: we first summarise the pedagogical rationale that suggests interactive engagement is an essential ingredient to encourage deep learning. We then consider practical issues of hardware, cost and installation before considering the pedagogical aspects of what constitutes a good question. We highlight a number of possible use scenarios for these systems and also what we have learned from an extensive evaluation of both students and staff experience.

There is already considerable activity with these systems within the UK. Our focus here will be on the application within the Physical Sciences. Three invaluable resources stand out as information goldmines. The first is the online material maintained by Steve Draper from the Department of Psychology at the University of Glasgow⁴. The second is the JISMAIL mailing list on electronic voting systems⁵. The third is the collection of resources at EDUCAUSE on the utilisation of ICT in higher education in the US, in particular the research bulletin devoted to transforming student learning with classroom communication systems⁶.

Pedagogy: interactivity as the essential ingredient

"The complex cognitive skills required to understand Physics cannot be developed by listening to lectures any more than one can learn to play tennis by watching tennis matches."
Hestenes⁷

Few academic practitioners would quibble with the notion of wanting to foster an environment where student learning is 'deep' rather than 'surface'⁸, enabling them to construct meaning rather than merely memorise facts. One characteristic of deep-learners are well-developed 'problem-solving skills', equipping the deep learner with the ability to tackle unseen problems beyond the confines of the presentation of the original material. McDermott⁹ has termed it "meaningful learning", connoting the ability "to interpret and use knowledge in situations different from those in which it was initially acquired". Development of such higher level skills is also one of the most difficult elements to 'teach'. Student activity offers a pathway to promote the processes of deep learning and develop student proficiencies in doing it spontaneously. In active learning (which has been termed "interactive engagement" by the Physics education research community) the student acts out the higher level cognitive processes of questioning, reasoning, organising and integrating within the subject context. The inclusion of peers in the process, through discussion, generates *inter*-activity. On a superficial level, such interactivity can address the attention span limit and can make the lecture a more enjoyable experience for students. On a more substantive level, engagement with the material and its underlying concepts has been shown to have a profoundly positive effect on student learning.

Hake¹⁰ has presented results of a six thousand student survey, assessing the efficacy of (differing elements of) interactive engagement in teaching as compared to more 'traditional' methods of instruction. The testing instrument has generally been one or both of the two diagnostic tests developed in the US as a measure of proficiency of understanding of fundamental concepts in mechanics: the Force Concept Inventory (FCI)¹¹ or Mechanics Baseline Test (MBT)¹². His bottom line conclusion from a statistical analysis

of the data is that the use of interactive engagement strategies "can increase mechanics course effectiveness well beyond that obtained with traditional methods". Though much of the work has been in the area of Physics, it seems that the applicability of these conclusions is not limited to this discipline^{13,14}, but can have an impact across many courses with challenging concepts. A recent paper in the domain of computer science has echoed these findings¹⁵. In addition, the wide-ranging list of subject areas this methodology is currently being applied to is further evidence, however anecdotal, of its effectiveness⁴. Use of these methods provides important feedback to all concerned. For staff, it enables the cohort's collective understanding to be gauged and for the students it allows formative assessment of their own progress.

Logistics: hardware, cost and siting

"Electronic voting systems typically comprise four elements: a tool for presenting lecture content and questions, electronic handsets, receivers that capture student responses and software that collates and presents students' responses"
Kennedy and Cutts¹⁵

There is now a bewildering array of vendors who can supply hardware, software and handsets (see Draper's pages for an up to date list¹⁶ and this survey for use in secondary schools¹⁷). A critical decision to be taken relates to the way the handsets transmit to the receivers; infra-red handsets generally cost much less (at least half, possibly a third the price) than those using radio-frequency communications. The downside to the infra-red hardware is that they are less reliable, need a receiver per 50 or so handsets, and these receivers must be carefully positioned around the lecture theatre to maximise the opportunities to collect all the votes in as short a time as possible.

In Edinburgh, the large class sizes determined that we bought the cheaper of the two alternatives; an IR-based system from GTOcalComp known as PRS (Personal Response System)¹⁸. In the summer of 2005, the cost of hardware (12 receivers, 400 handsets, adapters and brackets) was approximately £14,000. (This is to be compared to an approximate cost of £21,000 for the same number of radio frequency handsets.) We evaluated two different handsets on trial; the PRS ones were not ideal as there was no clear signal *on the handset* that the student had voted. On the alternative handset that was tested, there was a light to indicate that a vote had been successfully cast. However this handset also looked and felt far less robust than those of the PRS design. As always, these choices amount to a compromise and the need to match the educational requirements to the capability of the system. Our method of use was restricted to single-vote answers to multiple choice questions (MCQs), which we address shortly. There are far more sophisticated handsets (currently all RF) allowing, for example, text entry. There is the added steady-state running cost to be included, which we estimate at approximately 5% for lost or broken hardware, batteries etc.

Our IR handsets have come to be colloquially called 'clickers' (a nickname that originated in the US, and more than once has resulted in the confusion to prospective adopters that they 'actually click!') IR clickers must have a clear line of sight to a single receiver; signals cannot pass through desks or the heads of people in front of you. This dictates that the receivers (we have used 4 in series for a class of 250, 7 for a

class of 350) are mounted in an elevated position, well-separated from each other. In the theatre that accommodates 250, we have placed two at the front of the class, one either side of the teaching wall and two halfway up the lecture theatre, one on either wall. We instruct students to aim for the one closest to them, even though that might be (in true airline-safety-briefing style) behind them.

All systems come with software to collate and display student votes, some (eg the PRS software) with a plug-in for Microsoft *PowerPoint* that enables questions to be embedded within a slideshow and automatically started. We have found one needs the entire display screen to project a response grid which enables the students to identify that their vote has been received. The display of the question on which the students are voting, which clearly must be visible during thinking time, necessitates the use of second screen, overhead projector, or board.

The logistics of providing the students with handsets must be considered. We issue handsets at the beginning of the course and collect at the end, thereby avoiding the loss of valuable lecture time with distribution and collection of handsets. As the adoption of this as a lecturing technique becomes more widespread across other Schools in the College (akin to Departments within Faculties), we are investigating a centralised service of dispensing and collecting the handsets. We refrain from detailing the exact mechanics of operation

of an electronic voting episode within a lecture (several clear accounts exist elsewhere¹⁹, including a JISC-produced Innovative Practice case study video²⁰ and an EDUCAUSE podcast²¹). Photographs of the handsets and the lecture theatre set-up are shown in Figure 1.

Pedagogy again: what makes a good question?

“Although multiple choice questions may seem limiting, they can be surprisingly good at generating the desired student engagement and guiding student thinking. They work particularly well if the possible answers embody common confusions or difficult ideas.”

Wieman and Perkins²²

We have exclusively used multiple choice questions (MCQs) as interactive engagement exercises within our lectures in Edinburgh. MCQs have their supporters and opponents, but



Figure 1: Photographs of the handsets used in Edinburgh and a lecture theatre in use

for us this was a matter of practicality. We have accumulated (over a period of several years) a bank of some 400 MCQs relating to a first year Physics course in the classical study of space and time. In fact, in previous years, we have operated a low-tech version of the interactive episodes in lectures in which students used three coloured cards to indicate their response to MCQs. There are many reasons why the

electronic system is better (see for example the student quote within the 'Evaluation' section) but using the coloured cards system over time has provided us valuable insight into what it is that makes a 'good question'.

excellent example that evidences such misconceptions is illustrated in Figure 2 (actually taken from a diagnostic test given to entrant students at Edinburgh). Not only did the majority of the students who answered the questions answer incorrectly, they all chose the same incorrect answer. Such

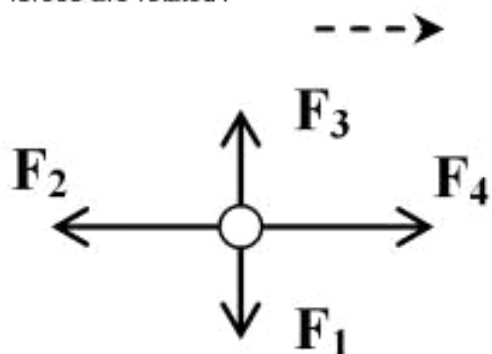
unanimous misunderstanding is very rare and in this instance can be traced back to elucidate the misconception that guides student choice: in this case the pre-Newtonian view of 'motion-implies-a-force'.

Our questions tend not to be overly numerical; any mathematics required is at the level of mental arithmetic only. Instead, they focus on concepts and the testing of the understanding of such concepts. The Physics Education Research (PER) community in the US has a long tradition of articulating the key requirements of and developing excellent MCQs for use in Physics. Developments have stemmed from Mazur's book a decade ago²³ describing a methodology known as Peer Instruction (to which we return later), to the Project Galileo website that Mazur created to collect many of these questions²⁴, to recent reports developing and extending Mazur's ideas²⁵. The widespread adoption of Mazur's approach has percolated upwards from classical and introductory mechanics and kinematics²⁶ into more advanced topics in Physics²⁷ and sideways into Chemistry²⁸.

Closely allied to the question of 'What makes a good question?', and already hinted at in the previous paragraph, is 'Where do I get good questions from?' It is undoubtedly true that many people harbour collections of such questions locally.

The annual marking of exam scripts is a good hunting ground for persistent misconceptions. Nearly all undergraduate text books now come with end-of-chapter questions of this type, many offering the full online provision of formative and

Four forces, F_1 , F_2 , F_3 and F_4 are exerted together on a hockey puck. The puck moves at constant speed along a straight line in the direction of F_4 . The arrows in the accompanying figure represent the directions of the four forces but not their magnitudes. Which of the following relationships represents best how the magnitudes of the four forces are related?



- a. $F_4 = F_2$ and $F_3 = F_1$.
- b. $F_4 = F_2$ and $F_3 > F_1$.
- c. $F_4 > F_2$ and $F_3 > F_1$.
- d. $F_4 > F_2$ and $F_3 = F_1$.
- e. $F_4 > F_2$ and $F_3 < F_1$.

Response Summary

Answer	Value	Frequency Distribution
-	0%	5
a	100%	13
b	0%	0
c	0%	0
d	0%	56
e	0%	0

Figure 2: Sample question and response profile, illustrating the overwhelming misconceptions that persist about classical descriptions of motion and forces

A good question is one where a spread of answers might be expected or where it is known that common misconceptions lurk. A poor question, by contrast, might be a deliberate trick question, or one that is distracting to the material at hand. An

summative testing mechanisms²⁹. This is a clear example where the growing digital object economy can prevent reinvention of wheels; the challenges are to be able to (a) discover resources and (b) interoperate different resource formats where necessary. Substantial local activity has been undertaken in the Physical Sciences Centre recently with this in mind; our own development project³⁰ to produce a browsable library of MCQs for selective download in different output formats; and the Centre's QuestionBank project aim to facilitate storage, sharing and interoperability of these (and other) resources.

Use scenarios

"Electronic classroom response systems...are merely tools, not a 'magic bullet'. To significantly impact student learning (they) must be employed with skill in the service of a sound, coherent pedagogy. This is not easy."
Beatty et al.²⁵

We have employed these interactive question episodes in a variety of different ways throughout our first-year Physics and Biology courses and we present some of these as potential use scenarios, which illustrate a number (but by no means an exhaustive list) of the ways that they can be used during lectures. The four scenarios that we outline represent progressively increasing departures from the traditional lecture format and similarly, can be thought of as progressively more challenging to introduce from the perspective of the lecturing team. We believe there is good justification for a mixed mode of use, a point we return to in more detail in the following section. We make no reference here to more complex uses of the technology, for example the handsets that allow transmission of text answers rather than simple numerical choices. Nor do we report the use of these techniques to serve as preparation for examinations³¹.

In its simplest form, these questions may be employed to simply break-up the lecture, to regain audience focus and attention and as a mild diversion to the main business of the lecture time around halfway through. Questions need only be loosely coupled to the course material (if at all), and are particularly suited to ice-breakers for a class new to this methodology in lectures. Sample questions include polling the age profile of the class, the subject study level at school, what the class has found most difficult / enjoyable etc. Clearly such questions cannot be used too frequently, otherwise the class will perceive them of little value and engagement with the process will wane. They are, however, particularly useful in familiarising the students and staff in the use and operation of the system.

A slightly more sophisticated and beneficial use of these episodes is to serve as a refresher or test of understanding of key points from material previously covered. An example would be to use a question at the beginning of a lecture, addressing material covered at the last lecture, potentially several days ago. This serves the dual purpose of 'heads-on' engagement of the student with material from the start of the lecture, and as a tool for the lecturer to gauge where the class is with regard to recall or understanding of previous key

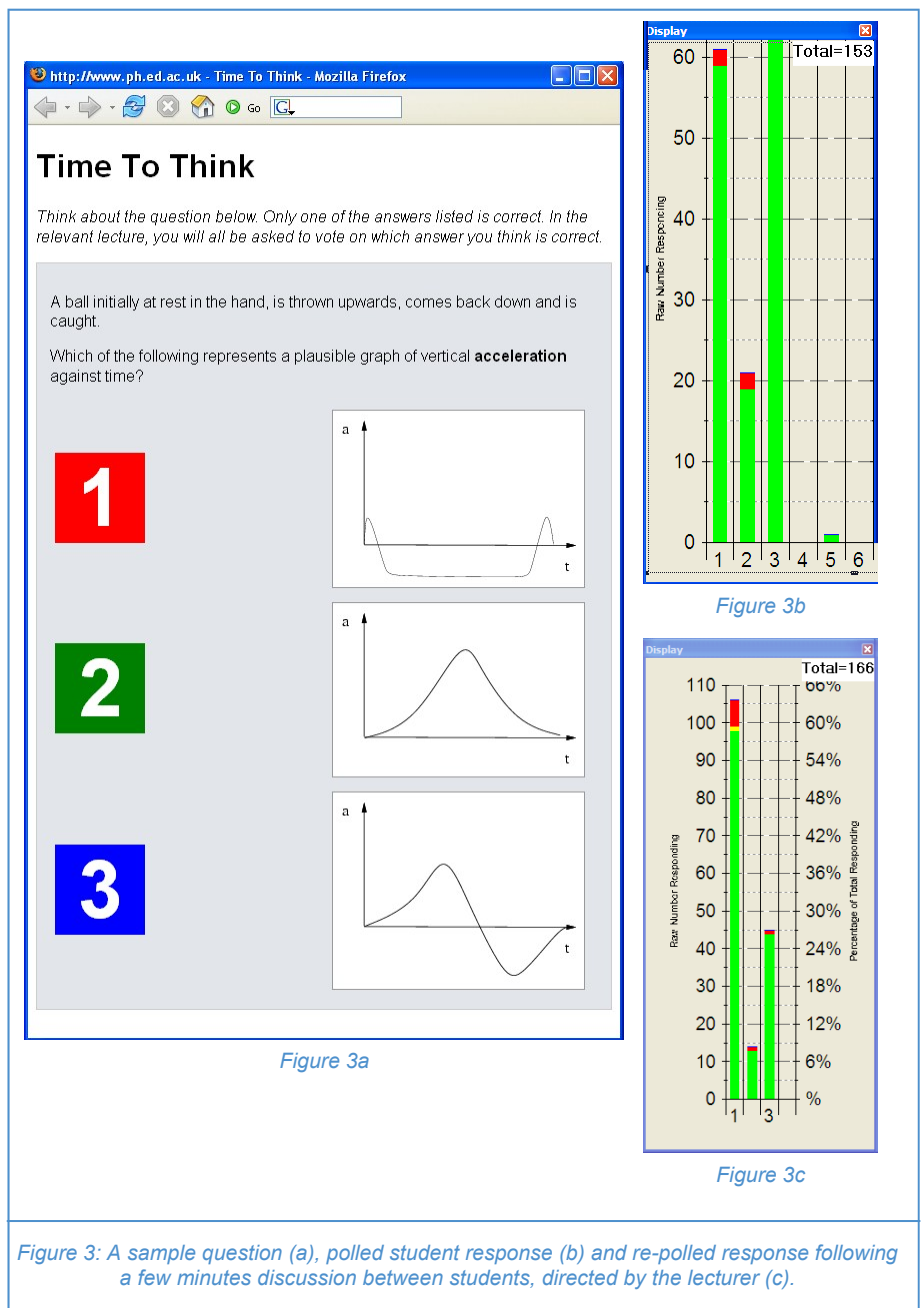


Figure 3: A sample question (a), polled student response (b) and re-polled response following a few minutes discussion between students, directed by the lecturer (c).

material. Wit³² has described such an approach and the subsequent actions of the lecturer in the light of student responses.

Progressing further along the scale of complexity brings us to the use of these questions as a vehicle for peer instruction²³, capitalising on the social context of discussion and peer interaction. The process for one of these episodes is that a question is posed and voted on individually. Following display of the class responses, students are invited to talk to neighbours and defend or promote their own viewpoint and why they think it is correct. The class is then re-pollled and the revised response distribution is displayed. This approach has been extensively used by Jim Boyle in Mechanical Engineering at Strathclyde²⁰, for which the physical layout of the lecture theatre was altered to foster small-group

are shown in Figure 3c. There is a significant swing towards the correct answer. The misconception implied by answer 2 is retained for a small proportion of the class and subsequent discussion of the problem highlighted why the incorrect answers were wrong. This is, for the lecturer new to this methodology, a potentially disruptive experience; letting go of the control of the lecture for a couple of minutes free discussion by students. In addition, an episode of peer instruction such as this begins to occupy a non-negligible fraction of the lecture time (perhaps 10 or 15 minutes). This has knock-on consequences for coverage which we return to shortly.

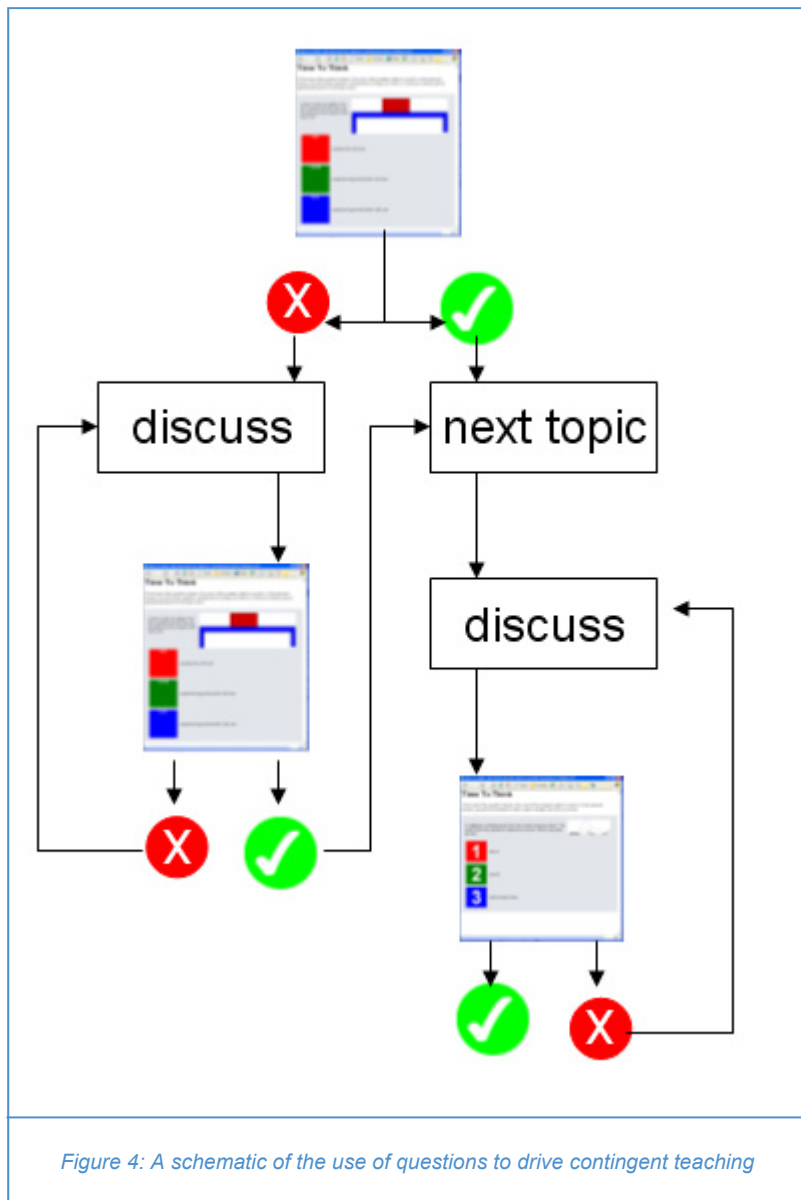


Figure 4: A schematic of the use of questions to drive contingent teaching

discussion, using crescent-shaped desks seating four students.

The results from our own experiences using this method are illustrated in Figure 3. The question (Figure 3a) is polled (answer 1 is correct) and the initial response histogram is shown in Figure 3b. The class was then invited to discuss responses with each other and re-pollled, the results of which

The final use scenario is what Steve Draper has termed “contingent teaching”¹⁹, where the interactive engagement episodes act as branch points in the lecture. Subsequent progression beyond these points is contingent on the response from the students. A question which, for example, 80% of the students get wrong would indicate either a fundamental misunderstanding associated with the material, or a lack of clarity in the exposition of it, or both. It would be negligent to pass over this without further comment and perhaps re-polling of the same or similar question. A schematic of how such contingencies might unfold is illustrated in Figure 4.

In this respect, the lecture truly becomes a two-way experience, a conversation between staff and students, mediated by the technology. The students collectively influence, perhaps even determine, the route that the lecture takes through the available pathways of material and questions. The thought of leading such sessions might be deeply uncomfortable to some staff and it is true that a great deal of care and planning (and limiting the number of possible contingencies) is needed to run these lectures effectively. It can be both an unsettling and, when it works, exhilarating experience to teach in this manner. Beatty et al.²⁵ have termed this method of instruction “Question-Driven Instruction” and present an explicit model detailing pedagogical goals and the mechanisms that achieve these.

The more complex these episodes within a lecture generally become, the more time they will occupy. It is important not to rush through them, but to give the class adequate thinking time (usually we choose about 2 minutes, but increase or decrease where appropriate). A cycle of peer instruction can take 10-15 minutes, perhaps longer if it is preceded by an orientation to the problem or topic at hand. This is a substantial fraction of the lecture slot and it raises issues of what to do about the content that would normally have been covered. Two equally unsatisfactory extreme views are (i) to try and shoehorn material in, or (ii) to not be concerned

about lack of coverage. We have found that if the lecture time is spent not covering the traditional A-Z of the course material, this material must exist elsewhere and be able to be integrated into the course. This may take place in other teaching activities: workshops³³ or tutorials, for example. Or it may be located in students’ self-study time, in the form of online resources, or a textbook. We have made use of the concept of a ‘learning contract’ with the students. The

agreement made at the start of the course is that if we do spend appreciable amounts of time on interactive engagement in lectures, their part of the bargain is that the displaced content is covered in other activities, be it other face-to-face teaching times, or in self study.

Evaluation

“Our most important piece of advice..... is to pay critical attention to what happens when you do it. Your students are your best teachers”

Beatty et al²⁵

One of the most difficult things to evaluate after using this methodology is the quantifiable effect it has on student learning. The absence (or perhaps ignorance?) of such data, and a reliance on more qualitative evaluative instruments unfamiliar to those with a research-based scientific background, makes it difficult to persuade detractors of the technique's efficacy. However, there are studies that attempt to correlate employment of this technique with assessment performance. The extensive study by Hake¹⁰ provides compelling evidence for the effectiveness of interactive engagement as a broad-ranging premise to improve student learning. Kennedy and Cutts¹⁵ have found a positive association between usage of an electronic voting system and learning outcomes tested in course assessment. El Rady¹⁴ has found a similar positive association. However, caution should be exercised before proclaiming this as a general finding. There are still many variables, the effects of which may be subtle and resolved only at the local level of course delivery, modes of use of these exercises, assessment methods etc. These may contribute to or even swamp any effects that can be isolated to arise solely as a result of using this methodology.

Our own investigations of the correlation between lecture attendance in a first year Physics class (more accurately 'participation', as evidenced by a recorded vote from a handset) with end-of-course examination performance has yielded a positive correlation, albeit rather weak ($R^2 = 0.18$). The scatter illustrates the inherent problems collapsing multi-dimensional data from an extremely heterogeneous cohort onto a one-dimensional scale, analogous to using a one-dimensional radial distribution function as a measure of three-dimensional liquid structure³⁴. Furthermore it is not clear if performance in the exam is related to lecture participation via cause or effect.

We have extensively evaluated the attitudinal aspects of the use of this methodology, from the perspectives of both students and staff. A dedicated section in the end-of-course student questionnaire, using a combination of Likert-scale and free text questions, has generated plentiful and informative feedback. This was supplemented by focus groups with small groups of students. In the Physics course that had previously utilised the same interactivity with coloured cards, the student response was almost unanimously positive, with the clickers often rated as one of the best things about the course. One difference between using clickers to coloured cards is the anonymity it confers on the student and the ease with which the whole-class profile is visible immediately following a question. The following quote, from a Physics student asked about how they felt when getting an answer wrong, exemplifies this:

“At least no-one knows you got it wrong. You can still sit there, you are still motivated, you are not like ‘Oh, everyone saw that I got it wrong, I’ll just crawl into my corner and die”

The general objective of increasing student involvement and engagement with the material was widely recognised as a positive feature by the students:

“The questions make you actually think about what the lecturer has just said, which helps ‘cause sometimes it just goes in one ear and out the other”

It would be misleading for us to conclude that this is has been an entirely successful venture without problems. In fact, deployment of this methodology across the three large courses (two in Physics, one in Biology) has not always been straightforward, but has always been useful. Based on firsthand reflections on the process, incorporating the feedback from other staff and students involved, we can offer some tentative ‘Do’s and Don’ts’ for people thinking of or actually taking their first steps down this road. Many of these suggestions arise directly as a result of our own (sometimes uncomfortable) experiences. There probably is no single universal recipe for guaranteed success, only pointers that can contribute to the likelihood of effective incorporation. Thus we suggest the following guidelines and advice:

Well before the lectures

- It should be recognised that incorporating this technique will take careful planning; do not under-estimate the time and effort taken. It is not an out-of-the-box solution.
- It is important to consider the requirements and impacts of the employment of this technique in all aspects of the design and operation of the course; it is not a simple add-on.
- Other staff involved in teaching the course must be on-board, adequately briefed as to what is trying to be achieved and trained in use of the hardware and software.
- Finding questions can be difficult. Finding or devising ‘really good’ questions can be very difficult, but examples are out there.
- Consideration of how many questions to use, where to place them and what modes of use to employ is important.

During the lectures

- The role of the first lecture is crucial: be clear about the ways in which you will use the system and what you expect of the students.
- Do have a ‘plan B’ in case things go wrong (they can!)
- Do have a ‘don’t know’ option for responses to minimise guesswork.
- Be prepared to increase or decrease the thinking time allocated to particular questions. The noise level within a theatre is a good indication of collective progress!
- Be agile enough to adapt things; perhaps allowing some discussion between students of a question that has illustrated more than one misconception.
- Do devote enough time to the question episodes, particularly to discussing why incorrect answers are wrong.

After the lectures

- Do reflect. Consider re-using particular problem questions in revision sessions or assessments to reinforce concepts.
- Evaluate the usefulness of certain questions; if the class all got it correct, you might want to revise that question for a future instance of the course.

Conclusions

"We want our students to be trained for life as independent learners because we know that, in the real world, few people spend any time learning much by means of lectures."

Raine³⁵

We have described the rationale behind and implementation of interactive engagement exercises in large class lectures, using an electronic voting system. This is an example of a 'disruptive technology' that challenges the role and utility of the traditional lecture in University education. It is a relatively widespread practice now and there are considerable resources available to help with matters of implementation and question design. We have found that wholly successful implementation is not always achieved in a course with no prior experience of this methodology and a pragmatic, incremental approach to its deployment is warranted, supplemented by staff reflection on the process.

The benefits to student learning and attitudes towards this are difficult to assess quantitatively. Yet there may be positive knock-on effects in terms of students' attendance at lectures (they are perceived as more useful, interesting), in student motivation and – particularly in the first year of study – the orientation to learning in a University context. The potential for change is encapsulated in the words of an anonymous first year student:

"I find I am even having to think in lectures"

Acknowledgements

We gratefully acknowledge the helpful advice provided during this work by Quintin Cutts and Steve Draper from the University of Glasgow. Financial support for the project, led by Paul McLaughlin of the School of Biological Sciences, was provided by the University of Edinburgh's e-Learning Project Fund.

References

1. National Committee of Enquiry into Higher Education: Dept. for Education (1997). Available online at <http://www.leeds.ac.uk/educol/ncihe/>
2. See Prensky's papers on this and the subsequent consequences for the way information is processed, eg <http://www.marcprensky.com/writing/>
3. Woodie Flowers, Professor of Mechanical Engineering from MIT, in an online lecture from 2000, available at <http://hitchcock.dlt.asu.edu/media2/cresmet/flowers>
4. Steve Draper's collection of pages for the Interactive Lecture Interest Group, <http://www.psy.gla.ac.uk/~steve/ilig>
5. Sign up at <http://www.jiscmail.ac.uk/lists/ELECTRONIC-VOTING-SYSTEMS.html>
6. <http://www.educause.edu/LibraryDetailPage/666?ID=ERB0403>
7. Hestenes, D., *Am. J. Phys.*, **66**, 465-7 (1998)
8. Entwistle, N. J. and Ramsden, P., *Understanding Student Learning*, Croon Helm, London (1983)
9. McDermott, L., *Am. J. Phys.* **61**, 295-298 (1993)
10. Hake, R. R., *Am. J. Phys.* **66**, 64-72 (1998)
11. Hestenes, D. and Wells, M., *Phys. Teach.* **30**, 141-158 (1992)
12. Hestenes, D., Wells, M., Swackhamer, G., *Phys. Teach.* **30**, 159-166 (1992)
13. Hake, R., *Conservation Ecology*, **5**, 28 (2002)
14. El Rady, J., *Innovate*, **2**(4) (2006). Available online at <http://www.innovateonline.info/index.php?view=article&id=171>
15. Kennedy, G. E. and Cutts, Q. I., *J. Computer Assisted Learning* **21**, 260-268 (2005)
16. <http://www.psy.gla.ac.uk/~steve/ilig/tech.html>
17. <http://www.abernet.org.uk/bigdownloads/Falkirk.pdf>
18. <http://www.gtcocalcomp.com/interwriteprs.htm>
19. <http://www.psy.gla.ac.uk/~steve/ilig/local.html>
20. www.elearning.ac.uk/innopract/practitioner/strathclyde.html
21. <http://connect.educause.edu/blog/dianao>
22. Wieman, C. and Perkins K., *Physics Today*, 36-42(2005)
23. Mazur, E., *Peer Instruction*, Prentice Hall, New Jersey (1997)
24. <http://galileo.harvard.edu>
25. Beatty, I. D., Gerace, W. J., Leonard, W. K. J., Dufresne, R.J., *Am. J. Phys.* **74**, 31 (2006)
26. See the collection of Peer Instruction questions maintained by Joe Redish at <http://www.physics.umd.edu/perg/role/PIProbs/ProbSubjs.htm>
27. Heron, P. R. L. and Meltzer, D. E., *Am. J. Phys.*, **73** 390-394 (2005) and refs therein
28. <http://www.jce.divched.org/JCEDLib/QBank/collection/ConceptTests/>
29. See for example <http://www.masteringphysics.com/>
30. <http://www.physsci.heacademy.ac.uk/Resources/DevelopmentProjectsSummary.aspx?id=218>
31. Draper, S. W. and Brown, M. I., *J. Computer Assisted Learning*, **20** 81-94 (2004)
32. Wit, E., *MSOR Connections* **3**, 5-11 (2003)
33. *CAL-laborate*, UniServe Science International Newsletter, 14 (2005) available online http://science.uniserve.edu.au/pubs/callab/vol14/cal14_bates.pdf
34. Further discussion of this is beyond the scope of this paper, but it will be presented at the JISC conference on Virtual Learning Environments <http://www.online-conference.net/vle2006/introduction.htm>
35. Raine, D., *Physics World* (December 2000) <http://physicsweb.org/articles/world/13/12/2/1>