Virtual Investigations : ways to accelerate experience

PROCEEDINGS

John Garratt

Department of Chemistry, University of York, Heslington, York YO1 5DD.

We need ways of teaching which will help to change student's approach to learning from one in which they perceive chemistry as being about 'finding right answers' to one in which they see it as a subject in which, although 'nothing is known absolutely', some important insights can be gained through investigations.

Two specific approaches aimed at achieving this goal are described. Both are aimed at complementing but not replacing laboratory work. They add an extra dimension to recipe laboratories by providing opportunities for students to make their own judgements about procedures.

One uses published papers as a framework within which students discuss how they would tackle the problem addressed in the paper, evaluate the methods actually used, and critically appraise the results obtained.

The second approach uses computer simulations of laboratory procedures. These can be used as a preparation for laboratory work, or to carry out a virtual investigation.

Introduction

I became interested in developing strategies for teaching and learning chemistry when I realised that students were not learning what I want them to learn. I want them to learn to think like a professional chemist. This involves, amongst other things, recognising the level of uncertainty in scientific knowledge, using knowledge creatively, and questioning (or critically evaluating) the authority of books and tutors. In contrast, many students seem to be content to be given 'the right answer', and to be told what they need to know.

I wanted to help them to acquire knowledge, and the skills to exploit it. I was failing. I began to ask what I could and should do about it.

In this paper I analyse the problem in more detail, and I describe two methods for helping students to develop skills needed by professional chemists which are specifically related to practical work. Subsequent papers in this volume[1,2] deal with the development of other professional skills.

The problem of right answers

Perry described a series of nine stages through which students progress[3]. These have been related specifically to the learning of chemistry by Finster[4,5]. Students in Perry's Position 1 see the world in dualistic fashion (right/wrong, good/bad, etc) This position is usefully paraphrased as

"*Right answers to everything exist. These are known to authority, whose role it is to teach them*".[6]

Subsequent stages involve gradual loss of confidence in this comfortable model of a world in which knowledge is certain. However, all being well, students eventually progress to the final stages where they recognise that knowledge is contextual and relative and that this 'relative' view has important implications.

Using Perry's model as a framework it was clear that my problem was that many, perhaps even most, students were still in Perry's Position 1. Furthermore, my style of teaching (and, I fear that of many other academics and of the text books they recommend) was more likely to reinforce this mental state than to change it. I was teaching *as though* there is a right answer, ***as though*** I know what that answer is, and *as though* the student's job is to learn it. Using Coldstream's words[7], I was coming dangerously close to '*colluding in a spoon-feeding process*'.

This clearer definition of the origin of my problem is only half of its complete analysis. The other half involves defining the attitude of mind which I wanted my students to develop. This is most useful if defined in clearly scientific terms. A seminal paper by Boothroyd[8] provides an excellent definition given by Feynman[9]:

"*Science is a way to teach how something gets to be known, what is NOT known, to what extent things are known (for nothing is known absolutely), how to handle doubt and uncertainty, what the rules of evidence are, how to think about things so that judgements can be made, how to distinguish truth from fraud, and from show.*" My problem is now defined more clearly. I need ways of teaching which change the students' approach to learning from one in which they perceive chemistry as being about 'finding right answers', to one in which they see it as a subject in which 'nothing is known absolutely' or, more generally, as Feynman saw science.

Problems without answers – an example

All chemists have an instinctive understanding of what they mean by a chemical reaction. However, at some (early) stage they become aware that some reactions do not go to completion. This introduces the concept of equilibria, and the concept that the extent to which a particular reaction will go is determined by its equilibrium constant. It is a small step to the idea that all reactions have an equilibrium constant - and therefore that no reaction goes to completion.

How does the student reconcile this idea with the common (textbook) statement that 'many reactions go to completion'?

One way to raise a discussion about this, which I have tried more often with groups of academics than with students, is

to ask for a 'rule of thumb' to use in deciding how big an equilibrium constant must be for a reaction to be regarded as one which 'goes to completion'. A common answer is that it needs to be bigger than 10^5 or 10^{10}, and a few people want it as high as 10^{20} or even more. Think about what these answers mean. If $K>10^5$ defines a reaction which goes to completion, the association of protons and hydroxide ions to form water falls within the rule of thumb, and the concentration of protons in water is negligible. If $K>10^{20}$ defines a reaction which goes to completion, HCl is not completely dissociated in water.

Of course this only demonstrates that 'it is a silly question', because the answer depends on the context. But it makes two serious points. One is that experienced chemists can, in their heads, convert between ΔG^θ and K, and can assess from the context whether it implies 'going to completion' or not; in particular they are aware that there can be special problems when the number of reactants and products are not equal. Thus we change our definition with the context. It is less clear that we explain to our students how and when we make these mental switches. The second serious point is that, even this elementary concept of equilibrium provides opportunities to *discuss* chemistry rather than just to *accept* everything as given. Overton[1] has more to say about this.

In the rest of this paper I will discuss ways of increasing the intellectual involvement of students with their practical work.

Practical work in chemistry

Chemistry is, without question, a practical subject and therefore laboratory work is an essential component of any chemistry course. But laboratory work is only part of the practical experience needed by a professional chemist. It cannot be too strongly emphasised that 'practical work' covers a much wider range of activities than is usually encountered in an undergraduate laboratory course, and it is mistaken to believe that laboratory work is likely to provide adequate practical experience for a career in chemistry.

Table 1 lists some aims of practical work. It is intended as an illustrative rather than a definitive list. The first four items on the list can be learned only through direct laboratory experience. Laboratory experience may (but often does not) provide a vehicle for learning the others.

A great deal of the laboratory work carried out by chemistry students, at least in the first two years of university courses in UK, involves following recipes[10]. The limitations of this kind of work has long been recognised[11]. Nevertheless recipe labs have the great advantage that they allow the inexperienced student to take the same attitude to laboratory work as is taken by the professional scientist: the recipe allows the student to devote all his or her attention to the *technique* and not to worry at all about *theory*. This point is illustrated in figure 1. The difference between a research programme and a student lab is not what goes on in the laboratory but in what goes on *outside it*. The defining, decision making, planning, interpreting is done by the research worker in the office, the library, the pub, the bath, or wherever seems suitable. Of course, similar stages are gone through for student laboratory exercises – but it is the tutor rather than the student who makes the decisions. Lack of involvement with the planning stages inevitably means that it is hard for students to see the need to do anything other than concentrate on producing high quality results and getting a high quality mark. This situation is not helped by putting distracting information into the lab manual. An instruction like 'add a two molar excess of ascorbic acid – how many grams is this?' is not the sort of question a well organised scientist asks in the laboratory; it is part of the preparation. As Johnstone[12] has pointed out, proper preparation or prelab work involves more than an instruction to 'read your manual before you come'.

A consequence of the lack of student involvement at the planning stage is that researchers and students working at the bench are likely to respond differently to the question 'what are you doing?' The former is likely to answer in terms of

Table 1: Some aims of practical work

Provide opportunities to develop:
- technical skill
- confidence in lab work
- observational skill
- awareness of safety
- recording skill
- data manipulation
- data interpretation
- presentational skills
- report writing
- oral communication

Provide experience of:
- designing an experiment
  - the experimental basis of theory
  - link between theory and practice
- consolidating subject knowledge
- the process of science

Figure 1: Schematic representation of the place of laboratory work in a research programme and a student practical course.

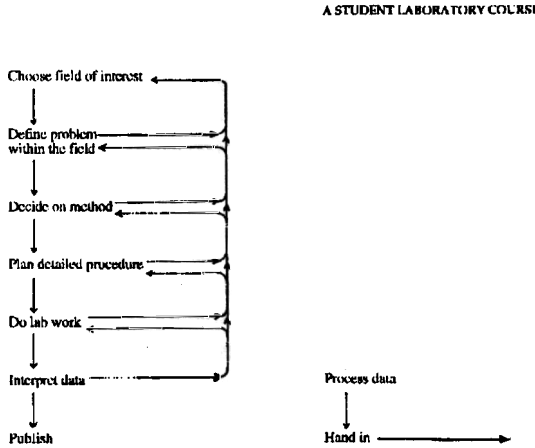

the relationship between the chosen method (stage 3 in fig 1) and the defined problem (stage 2 in fig 1). The latter is likely to say 'I'm adding that stuff, er, whatever it is, er some kind of acid, is it ascorbic?' or to point to their lab manual and say 'I'm here'.

The parallel with the kitchen recipe is precise. The inexperienced cook is not concerned with the thought processes of the *cordon bleu* chef who created the recipe, but only with cooking a dish to impress the guests at the dinner party. The student following the recipe is concerned only (or at best primarily) with getting a good result and hence a good mark; they are not concerned with fitting the experience of their laboratory work to their existing knowledge or with 'consolidating their learning by asking themselves what is going on in their own heads' [12]. We need teaching methods which involve the students in the stages of experimental design and data interpretation (which is more complex than the processing of data according to a given algorithm). Final year project work can provide these opportunities, but I suggest it is too little, too late.

I will describe two strategies for introducing some of these aspects of practical work into chemistry degree courses in the second, or even first, year of the course. These involve the use of published papers, and of computer simulations.

Scientific Papers

This approach is based on an original suggestion made by Brian Mattinson, was developed through discussions with him, and has been summarised previously[13,14]. The approach involves using a short published paper as a framework within which to discuss how the students would tackle the problem which is addressed by the paper. In York we developed four exercises, in conjunction with the author of each of the four papers which form the basis of the exercises[15-18]; these are now included in the Communications in Chemistry package available from the Royal Society of Chemistry[19]. We use these with classes of 25-35 students working in groups of four during a 3 hour class. I will illustrate the kind of tasks we set the students with some from a paper which describes the inhibition of the enzyme chymotrypsin by iodinated esters of tyrosine [17]. The complete paper is two pages in length.

The groups of students are first given the introduction to the paper, a small amount of background information, and 8 questions designed to set the work in context so that the students appreciate how the investigation fits in with other knowledge. The next 3 questions, related to measuring the rate at which chymotrypsin hydrolyses N-acetyltyrosyl-ethylester, are shown in figure 2.

The tutor in charge of the class allows 5-10 mins for groups to discuss these questions, and then collects answers for general class discussion. A positive feature of this approach is that it is hardly ever necessary to diminish student confidence or enthusiasm by telling them their answer is incorrect; no doubt it really would be possible to devise a way to measure the rate of reaction using sophisticated and expensive techniques such as nmr spectroscopy or mass spectroscopy! Particularly revealing are answers to the question about factors to be taken into account when choosing a method. Most groups think of factors such as sensitivity, accuracy, precision; but fewer think of availability of equipment or of cost, and usually only one or two recognise that familiarity with a technique ('I've done it before') is an important factor for many professional scientists.

During the discussion, students quickly recognise that most of the suggestions, though possible, are not practicable, and almost always discussion eliminates all but two methods (both

Figure 2: Part of the student handout for section 1 of a scientific paper exercise.

ESTERS OF IODINATED TYROSINE AS INHIBITORS OF CHYMOTRYPSIN

We have shown that chymotrypsin will not hydrolyse either 3-iodo-N-acetyltyrosyl ethyl ester or 3,5-diiodo-N-acetyltyrosyl ethyl ester and that these two compounds inhibit the chymotryptic hydrolysis of N-acetyltyrosyl ethyl ester.

The hydrolysis of N-acetyltryrosyl ethyl ester (ATEE) is shown below

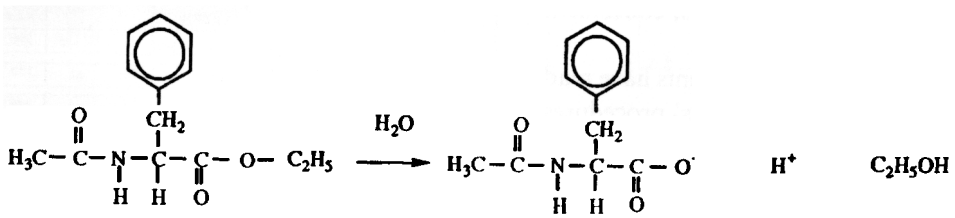

Now imagine that you want to conduct a similar study.

9. What methods can you think of by which you could measure the rate of formation of either product, or the rate of disappearance of the substrate?

10. What factors are likely to affect your choice of the most appropriate method?

11. Which of the methods do you think would be most suitable?

Figure 3: Fragment of the student handout for section 2 of a scientific exercise.

METHOD

The reaction was started by the rapid addition of chymotrypsin solution to the unbuffered substrate solution. The rate of reaction was followed by adding 10 mM NaOH from a microsyringe at the rate required to maintain the pH constant.

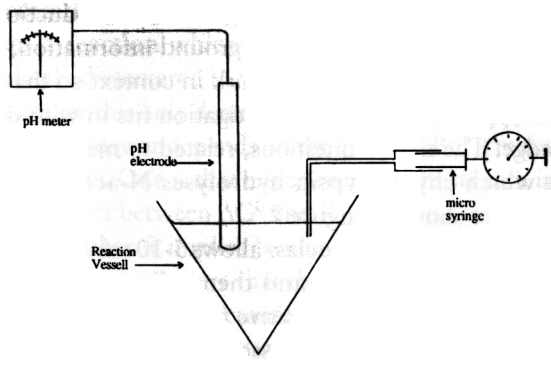

Now plan how you would carry out the experiment.

3. (i) In what kind and size of vessel will you put your ester solution?

   (ii) What would you regard as a convenient volume of solution to use?

Note: these are two of 6 questions related to the procedure.

of which have been published). One involves measuring the rate at which base must be added to an unbuffered solution in order to maintain the pH constant, the other involves measuring the change in ultraviolet absorption in a buffered solution.

The students are then given part of the methods section of the paper which states that the first of these two methods was chosen. This is the cue to asking them to discuss in some detail how they would set up the equipment to measure rate. Two of the questions they discuss are given in figure 3.

Of course there are no correct answers to these questions. My own view is that a beaker is the only reasonable reaction vessel, and that volumes outside the range 10-100 cm^3 are unreasonable. But this part of the exercise always stimulates fruitful discussion about the factors which affect the convenience of methods (manageable volumes, convenient times of measurement, etc).

By the end of the three hour class, the students have read almost all of the paper, compared the authors' procedures with the ones they proposed, and examined the results critically in order to assess whether the authors' conclusions are convincing.

I hope this brief sketch of the way we use scientific papers shows that they can provide students with opportunities to discuss and explore aspects of the design of an investigation and the processing and interpretation of data. The limitation is that the students can only compare their ideas with the published procedures; they cannot try out their own ideas. I will now discuss how far computer simulations can overcome some of these limitations.

The *eLABorate* project

The *eLABorate* project is concerned with the creation of computer simulations which provide students of chemistry and biochemistry with opportunities to design 'virtual investigations', and to process and interpret the resulting data[20, 21].

Simulations can support different aspects of learning. Our view is that they fulfil four roles particularly well. These are:

- as a preparation for laboratory work (eg *tracerLAB*);

Figure 4: Screen displays from *tracerLAB*.

(a) Selection of the amount of radioactive tracer.

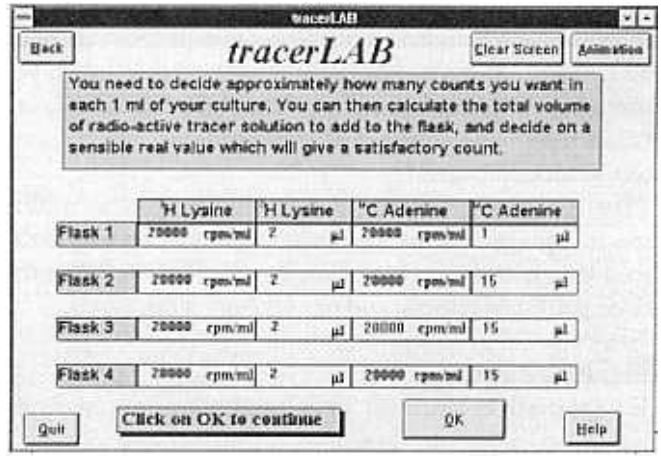

The user enters values in columns 2 and 4 to indicate the intended concentration of radioactive tracer in the incubation medium. Based on information on the specific activity of the stock solution of tracer and the expected efficiency of counting, the user calculates how much tracer to add to the incubation medium (columns 3 and 5).

(b) Selection of the amount of non-radioactive carrier.

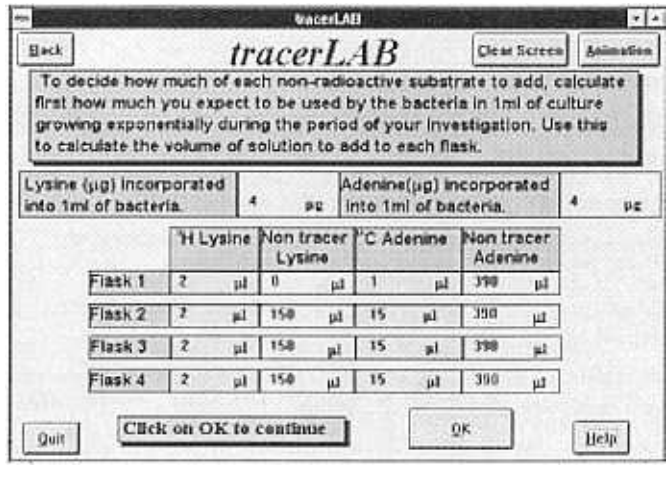

The user has already selected the total volume of growing bacteria and the volume of tracers (columns 1,2,4). Based on information about the density of bacteria in the medium and the expected growth rate, the user calculates how much tracer will be incorporated into protein and nucleic acid during the course of the experiment (boxes) and hence the amount of non-radioactive carrier needed (columns 3 and 5).

- to explore theory (eg *electrochemLAB*);
- to gain experience with expensive equipment (eg *nmrLAB*);
- to carry out a virtual investigation (eg *enzymeLAB*).

I will illustrate each of these uses of simulations with a brief description of the chosen examples.

tracerLAB

This programme simulates an experiment carried out at York by first year students of biology and biochemistry. The experiment involves the growth of bacteria in the presence of radioactive tracers (^{3}H-lysine and ^{14}C-adenine). These are incorporated into protein and nucleic acid. In the experiment, the students grow their chosen strain of bacteria in a medium containing one or both tracers, and follow incorporation of radioactivity into protein and nucleic acid. They do this during normal growth and after the inhibition of either protein synthesis or nucleic acid synthesis. Amongst the decisions they need to take in planning this experiment are:

- the amount of radioactivity needed to give a measurable count rate even in the early stages of the experiment;
- the amount of carrier (non-radioactive) material to add; this must be sufficient to ensure that all the radioactive substrate is not used up before the end of the experiment;
- the timing of the samples.

tracerLAB was written to allow them to plan this experiment before starting work in the laboratory.

Figure 4a,b shows selected screen dumps from *tracerLAB* illustrating how the students are guided to take these decisions. Figure 5a,b shows results from different simulated experiments. These illustrate some of the poor decisions which can be taken and which make the data uninterpretable.

Given the time it would take a student to obtain data for the graphs shown in figure 5, and the cost of doing so, it is not surprising that most organisers of laboratory courses provide a recipe rather than risk the students carrying out a badly designed experiment. However, the computer simulation can generate the data in a negligible time; and so students can plan their experiment, collect data, and revise their plans several times during a half day period. This gives them an understanding of the range of conditions which will provide useful results and hence, even if they carry out the actual experiment according to the tutors' protocol and not their own, they appreciate that the conditions were chosen with care and reasoning.

electrochemLAB

This program was written with the intention of helping first year students to develop a deeper understanding of the difficult concepts of basic electrochemistry. The package consists of six modules, each of which is presented to the students as a single screen. The modules deal with

- Redox Equilibria
- Qualitative Electrochemistry
- Quantitative Electrochemistry
- Activity
- Temperature Dependence
- Potentiometric Curves

Figure 5: Results from *tracerLAB*

Some poor judgements which can be made when planning an experiment with tracers. The results were obtained by setting up four flasks each containing ^{3}H-lysine and ^{14}C-adenine. Not all students recognise that they can distinguish between the two isotopes, and so use separate flasks for each type.

Data in (a) were obtained after removing the step which generates experimental error. This feature is not available to students. In (b) realistic experimental error has been reinstated.

(a) Incorporation of ^{3}H-lysine into protein.

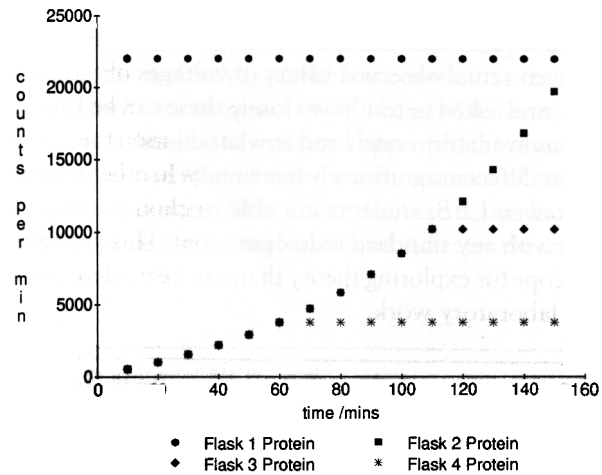

Flask 1: no carrier lysine added, so that all the radioactive tracer is used up in the first few minutes.

Flask 2: exponential incorporation for 150 min; the lower count rate at 150 min (compared with flask 1) shows that the carrier will last only a little longer than the experiment; but see (b) for effect of adding slightly less carrier than is ideal.

Flask 3: inhibition of protein synthesis at 100 min.

Flask 4: inhibition of protein synthesis at 50 min.

(b) Incorporation of ^{14}C-adenine into nucleic acid.

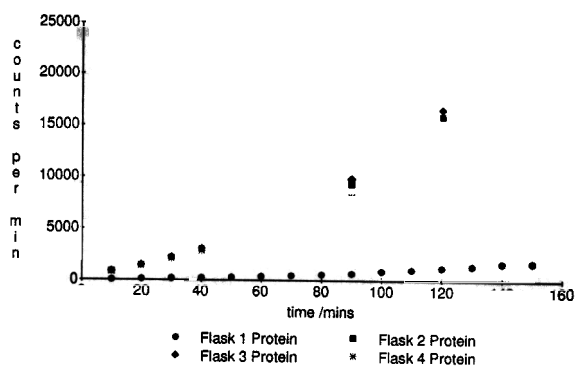

Flask insufficient tracer added; count rate in early samples is too low to be reliable.

Flask 2: exponential rate of incorporation for 140 min, showingthat ^{12}C-adenine is used up between 140 and 150 mins.

Flask 3: inhibition of protein synthesis at 100 min; incorporation of ^{14}C-adenine continues linearly, but inhibition is initiated so late that this is not clear; compare (c).

Flask 4: inhibition of protein synthesis at 50 min.

The fourth module dealing with activity coefficients illustrates the approach. The screen is shown in figure 6. As it shows, the user can set up an electrochemical cell in which one half cell is an iron II/iron III couple and the other is a standard hydrogen electrode. The concentrations of iron II and iron III and the charge on the anion are selected by the user. The program calculates and displays the voltage generated by the cell according to three models –

- an activity coefficient of unity for all constituents,
- the Debye-Huckel limiting law,
- the extended Debye-Huckel law involving the factor B.

There are many ways in which students can use this module to explore the concept of activity. For example, they can be given actual observed values of voltages obtained for real cells, and asked to test how closely these can be fitted by the various available models and at what dilution the activity coefficient differs insignificantly from unity. In other modules of *electrochemLAB*, students are able to choose imaginary half cells with any standard redox potential. This provides a greater scope for exploring theory than can be made available through laboratory work.

Figure 6: *electrochemLAB* : module 4 Activity coefficients

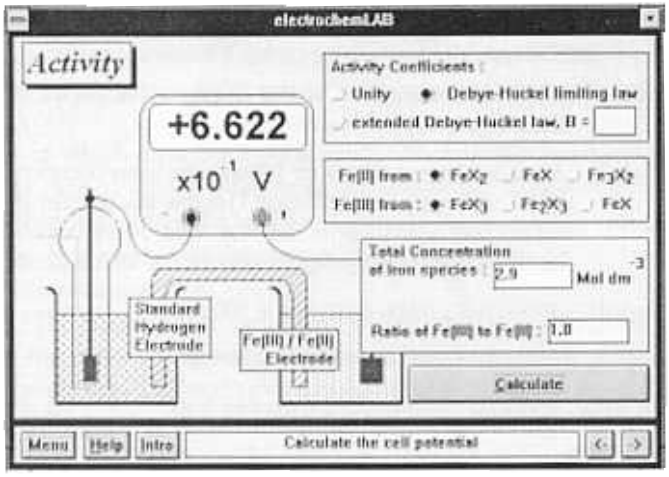

The user selects

- the method for calculating activity co-efficients;
- the charge on the anion in the FeIII/FeII electrode;
- the concentrations of FeIII and FeII.

The program then calculates the expected voltage.

nmrLAB

nmr spectroscopy is one of the most widely used techniques in molecular analysis. However, modern nmr spectrometers are clearly much too expensive for undergraduate students to be given access to them to run spectra. It follows that student experience with the technique is limited to basic theory and to the interpretation of spectra.

nmrLAB is designed to provide students with experience of the processes of data collection and manipulation as well as with data interpretation. Our expectation is that this will help them to make more effective links between theory and practice.

Figure 7: nmr LAB : the control of the spectrometer

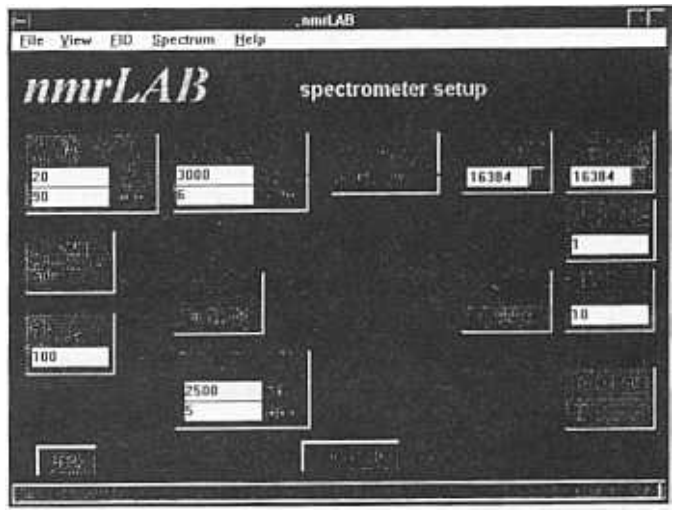

Figures in the white boxes are selected by the user.
Figures in the grey boxes are calculated from the selected values.

Figure 8: FIDs generated by nmr LAB

(a) based on 10 scans

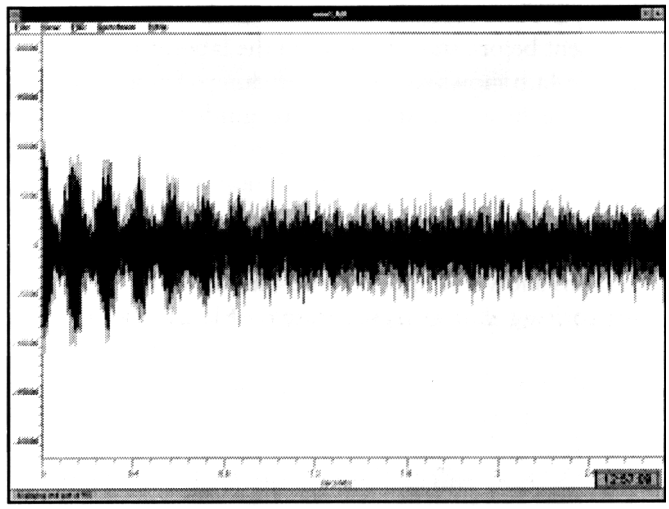

(b) based on 10,000 scans

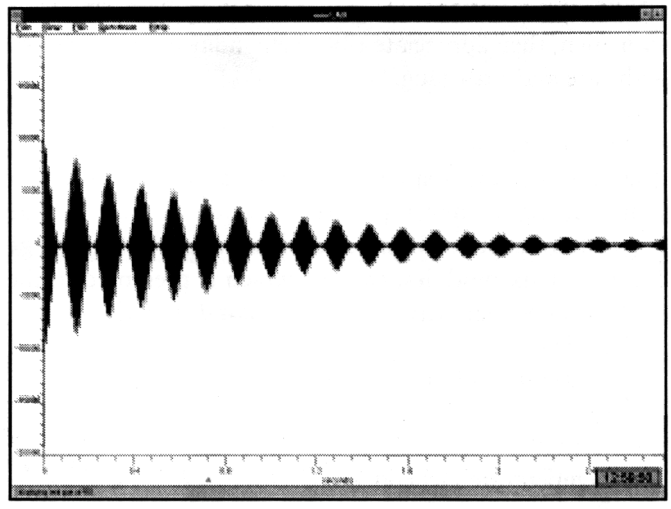

Figure 7 shows the initial screen of *nmrLAB* which represents the controls of an FT nmr spectrometer. The user can select all the settings available to a real operator – for example the instrument gain, number of transient data points, number of scans etc. The computer then simulates a FID for ethylethanoate based on the chosen settings. (Eventually we intend to include a library of compounds in addition to ethylethanoate, but this example illustrates the power of the approach). Figure 8 shows FIDs based on 10 and 10,000 scans; both are simulated in seconds even though a real FID based on 10,000 scans would take about 24 hours to collect. Figure 9 shows the resultant spectra.

Clearly it would take weeks, months or years to explore the effects of varying instrument settings on a real instrument in this way.

enzymeLAB

This simulation is based on a program developed some years ago to provide students of biochemistry at York with the opportunity to plan a kinetic investigation of an enzyme. The enzyme obeys Michaelis Menten kinetics – the rate of the enzyme catalysed reaction (v) is given by the expression

$$v = \frac{k_{cat} \cdot [E] \cdot [S]}{K_M + [S]}$$

In this expression k_{cat} and K_M are characteristic of the enzyme and the pH of the solution in which the reaction takes place. The values of k_{cat} and K_M and their pH-dependence are selected by the program for each user, and so each user is presented with a different enzyme which is 'realistic but not real'. [E] and [S] are the concentrations of enzyme and substrate respectively, and are under the control of the user. k_{cat}·[E] is defined as V_{max} (which is constant at constant pH and constant concentration of [E]). Figure 10 shows the main screen after the user has selected appropriate values and the program has simulated three measurements of rate under identical conditions. This illustrates that the measurements are subject to experimental error, and indicates the magnitude of the error.

Figure 9: Spectra generated by *nmrLAB* from the FIDs in figure 8.

(a) based on 10 scans

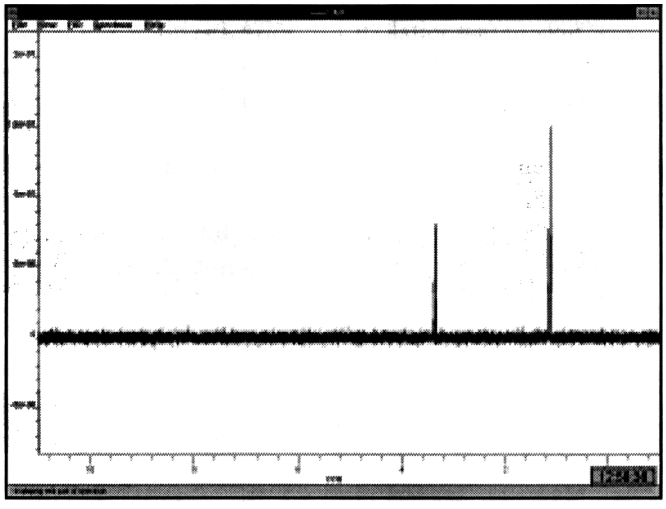

(b) based on 10,000 scans

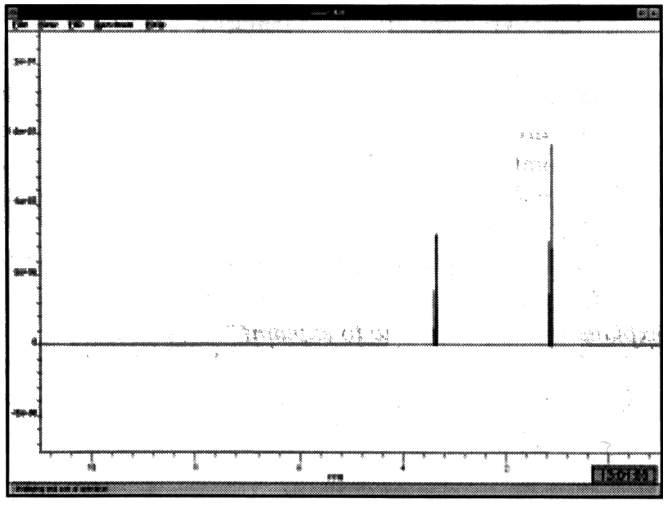

Figure 10: enzymeLAB: the main screen

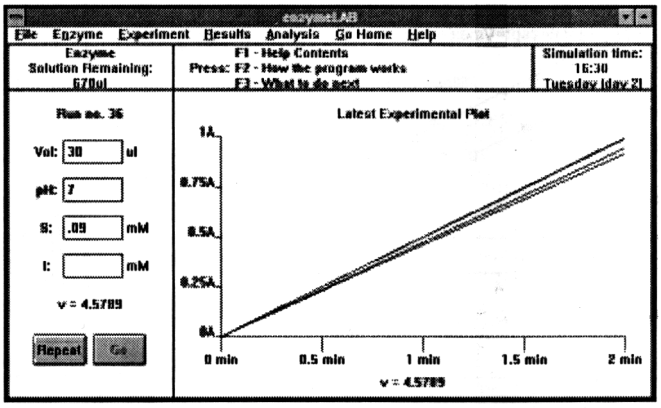

The user has chosen to use the repeat button to make 3 measurements of rate under identical conditions.

A variety of investigations can be carried out using this simulation. Our students are asked to find the optimum pH of their enzyme, V_{max} and K_M at this pH, and the effect of pH on K_M and V_{max}. It requires thought and skill to make appropriate choices of [S], [E], and pH at which to measure v in order to carry out this study effectively. Our students are required to report their results in two graphs. This requires them to plot different sets of data on a single graph. The plotting facilities within *enzymeLAB* do not provide a way to do this, because we want the students to make their own decision about how best to present their data. The program allows them to export their data to a spreadsheet with graph plotting facilities. Figures 11 and 12 show two graphs presented by different students. They demonstrate two valuable points. First, it would take a competent experimentalist several days to collect the amount of data shown, though only an hour or two to collect it via the

Figure 11: enzymeLAB : student results showing optimum pH; compare the presentation with that shown in figure 13.

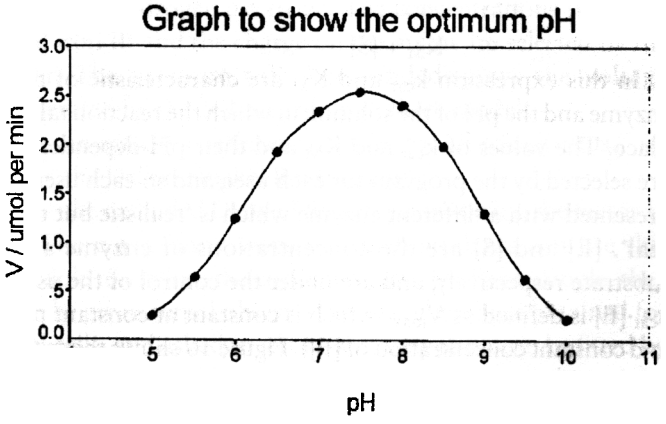

Figure 12: enzymeLAB: student results showing effect of pH on K_M and Vmax. Compare this with the alternative presentation in figure 14.

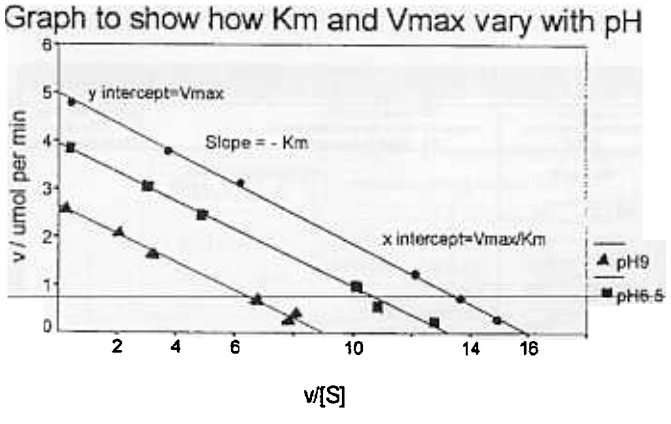

simulation. Second, the presentation of the graphs could be improved, showing that the students lack experience at presenting this amount of data so that this kind of exercise provides them with useful experience with data presentation. Figures 13 and 14 show equivalent graphs obtained and plotted for demonstration purposes by the *eLABorate* team.

Conclusion

The use of the scientific papers and of simulations, as described here, are not a good way of introducing students to new theory or knowledge. Rather, we see them as ways of involving students more closely with what Gott and Murphy[22] refer to as "*procedural understanding in science*". This procedural understanding is a very important aspect of science, and it must be learned through experience rather than taught didactically. Recognising this focuses our attention on the value of learning through experience.

To summarise what I have learned from my involvement with the teaching methods described here, I would like to use

Figure 13: pH profile of enzyme ref no: PR/57-240-010 Effect of pH on the rate of the enzyme catalysed reaction:rate is given in µmol of product per minute per ml of enzyme.

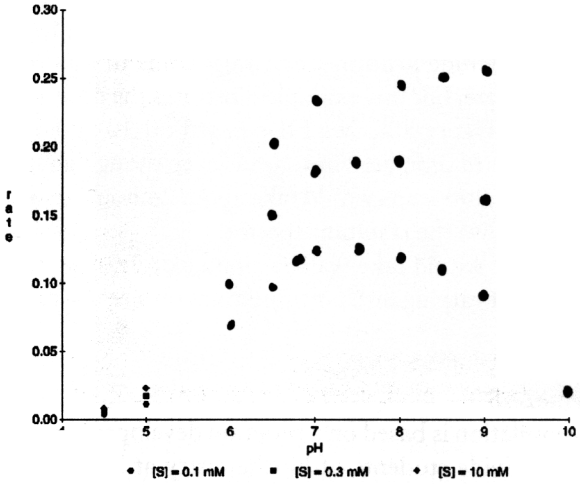

Figure 14: Eadie Hofstee plots at different pH for enzyme ref no: PR/57-240-010

slope = $-K_M$
y intercept = V_{max}

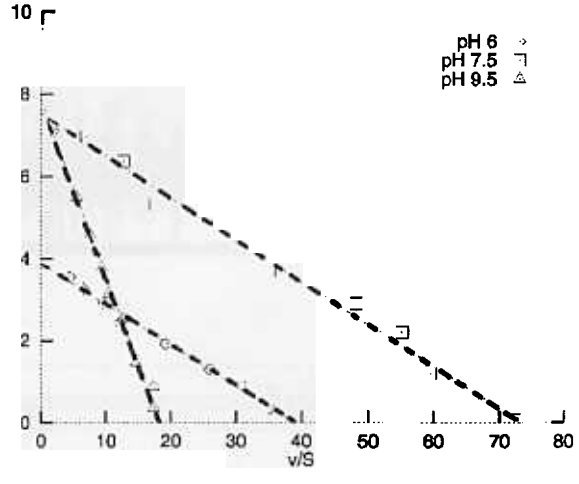

a quotation from Verdonk[23] and two propositions derived from Boothroyd[8].

My quotation is in two parts:

"Fact making has been replaced by fact learning by a process we call bookification."

"Should our students learn descriptions or learn how to describe; learn experiments or learn how to experiment; learn explanations or learn how to explain?"

Verdonk's concern that students need more involvement with 'fact making' or with the process of science leads directly to my two propositions which are:

1. We should talk more about learning and less about teaching; we should not confuse teaching with telling;

and we should remember the old adage that nothing is taught until something is learned.

2. When we plan our teaching (or, better, our students' learning) we should put less emphasis on the teaching of chemistry and more emphasis on learning how to be chemists; because being a chemist involves knowing chemistry, but knowing chemistry (alone) does not make a chemist.

Acknowledgments

None of this work is entirely my own. It is a pleasure to acknowledge the stimulating environment provided by the Department of Chemistry at the University of York, and in particular the contributions made by the many colleagues with whom I have been privileged to work. These include Andrew Horn, for helping me refine my example of 'problems without answers'; Brian Mattinson, who originally suggested how scientific papers could be used as a learning experience; and also Chris Anastasi, Pat Bailey, James Clark and Mary Aitken who worked with us to develop the four exercises now in use; the York *eLABorate* team of Doug Clow, Jane Tomlinson and Anne Hodgson whose creativity and imagination have been essential to the development of the concepts of simulations described here; those who identified needs for specific *eLABorate* simulations, guided us in their preparation, and used them with their students: Simon Hardy (*tracerLAB*), Paul Walton (*electrochemLAB*), Simon Duckett (nmrLAB).

The *eLABorate* project was funded by the University Funding Councils of England, Scotland, Wales and Northern Ireland under the Teaching and Learning Technology Programme. Additional financial support was provided by the Elaine and Margaret Moreton Trust of the University of York, the University of York Teaching Innovation Fund, and Norsk Hydro UK Ltd.

References

1. Overton T 1997 Creating critical chemists *U. Chem. Ed.* **1** 28
2. Bailey P 1997 Coaxing chemists to communicate *U. Chem. Ed.* 1 31
3. Perry WG 1979 *Forms of Intellectual and Ethical Development in College Years: a Scheme* (Holt, Rinehart and Winston, New York)
4. Finster DC 1989 Developmental instruction I *J. Chem. Ed.* 66 659-661
5. Finster DC 1991 Developmental instruction II *J. Chem. Ed.* 68 752-756
6. Phillips PS and Pennington M 1993 The Perry model for developmental instruction, in *Proceedings of Variety in Chemistry Teaching 1993* (ed. M Aitken) (University of York) pp 81-82
7. Coldstream P 1997 Chemistry education for a changing world *U. Chem. Ed.* 1 15
8. Boothroyd C 1993 Lets talk about learning, in *Proceedings of Variety in Chemistry Teaching 1993* (ed. M Aitken) (University of York) pp 55-65
9. Gleick J 1992 *Genius: Richard Feynman and modern physics* (Little, Brown and Company) 284-5
10. Meester AAM and Maskill R 1994 *Second year practical classes in undergraduate chemistry courses in England and Wales* (Royal Society of Chemistry)
11. Young JA 1957 Teaching the scientific method in college general chemistry *J. Chem. Ed.* 34 238-239
12. Johnstone AH 1997 Chemistry teaching – science or alchemy *J. Chem. Ed.* 74 262-268
13. Garratt J and Aitken M 1994 A new approach to the teaching of experimental design *Biochem. Ed.* 22 16
14. Garratt J and Overton T 1996 Scientific papers as a teaching aid *Educ. Chem.* 33 137-138
15. Anastasi C, Bemand P P and Smith IWM 1976 *Chem. Phys. Letts.* 37 370-371
16. Bailey PD 1987 *J. Chem. Res.* (S) 202-203
17. Garratt CJ and Harrison DM 1970 *Febs. Letts.* 11 17-18
18. Clark JH, Duke CVR, Brown SJ and Miller JM 1986 Infrared spectroscopic materials *Spectrochimica Acta* 42A 811-814
19. Bailey P and Shinton S *Communicating Chemistry* (Royal Society of Chemistry)
20. Garratt J, Clow D and Swift A 1994 The eLABorate project *Software Reviews* 9 14-15
21. Garratt J, Clow D, Hodgson A and Tomlinson J 1997 *UniServe Science News* 7 6-8
22. Gott R and Murphy P 1987 Assessing investigations at ages 13 and 15 *Assessment of Performance Unit Science Report for Teachers* 9 (DES/WO/DENI, London)
23. Verdonk A 1993 The role of educational research in the quest for quality in the teaching and learning of chemistry, in: *Proceedings of Variety in Chemistry Teaching 1993* (ed. M Aitken) 49-54 (University of York)