Computer-supported development of critical reasoning skills[1]

David Spurrett
University of KwaZulu-Natal, South Africa

ABSTRACT

Thinking skills are important and education is expected to develop them. Empirical results suggest that formal education makes a modest and largely indirect difference. This paper will describe the early stages of an ongoing curriculum initiative in the teaching of critical reasoning skills in the philosophy curriculum on the Howard College Campus of the University of KwaZulu-Natal (UKZN). The project is intended to make a significant contribution to the challenge of helping students think more effectively.

The general outlines of the critical reasoning skills 'problem' are described, along with some remarks on the form it takes in South African post-secondary education. This is accompanied with some observations concerning the optimistic claims often made on behalf of philosophy, and the study of philosophy, in the area of reasoning skills, and some indications of the actual success of most attempts to teach reasoning skills. Thereafter some general results from an approach in the study of cognition, most often referred to as 'distributed cognition', are outlined.

These results form part of the explicit motivation for the development and design of a software system for supporting critical reasoning teaching. The ReasonAble system was developed at the University of Melbourne, Australia, and is currently in use at UKZN. The main features of the software system are briefly described. Finally the specific implementation developed so far at UKZN is explained, and the results of initial evaluations by students are reported. Some comments on envisaged future evaluations and forthcoming initiatives extending the use of the system are outlined.

Keywords: Critical reasoning, curriculum, software, philosophy

INTRODUCTION

This paper describes the early stages of an ongoing curriculum development initiative in the teaching of critical reasoning skills in the philosophy curriculum on the Howard College Campus of the University of KwaZulu-Natal (UKZN). The project, intended to make a significant contribution to the challenge of helping students think more effectively, involves the use of Australian-developed software called *Reason!Able,* usually used in dedicated critical reasoning courses.

The software and related tutorial programme was introduced at UKZN in response to the regularly acknowledged 'problem' of students' critical reasoning skills. While the critical thinking problem is global, it is not evenly distributed, and it takes a particular form in South Africa. Remaining drastic inequalities in the education system mean that many South African students enter the post-secondary system relatively poorly prepared, while some are very well prepared. Levels of skill of various sorts, including critical reasoning skill, vary widely within single post-secondary classes, as do levels of motivation for critical reasoning.

Classes containing such variation in talent, skill and motivation demand creativity and dedication in the teaching process. There is a risk of sectors of any given class being left behind, driven

away, demoralised or sold short. The challenge is, perhaps, especially urgent at the early stages of undergraduate curricula, where rapid gains in skill and motivation at critical reasoning could help lay a foundation for future success in other courses and beyond. The project described here is partly responsive to this particular challenge – the urgent need to build skill and confidence in critical reasoning among students who enter the system often with little of either but a pressing need for both.

My involvement as a philosophy lecturer means that I work in a discipline, the study of which is *correlated* with highly developed reasoning skills. (In the case of philosophy, departments regularly claim that study of philosophy *causes* improvement in reasoning skills, but this has not been demonstrated.) There is also evidence that tertiary education does make a positive, but modest, difference to reasoning skill, and that typical reasoning courses make a comparatively small contribution to the process.

Some key premises about learning underpin the choice and use of the software in this project. For space reasons, the key constructs are summarised very briefly, while the references provided point to a more detailed explanation. In brief, the project is premised on related concepts from distributed cognition, which provide the inter-rated analytical categories of scaffolding, transformations, and the power of good representations.

Distributed cognition (Hutchins 1995; Clark 1997; Spurrett 2003) refers to cognitive processing that relies to a significant extent on resources outside the brain. Evidence from robotics, human computer interaction, developmental psychology, cognitive psychology, cognitive anthropology, as well as various parts of biology and other fields, is sufficiently strong to establish the claim that cognition is at least often distributed. This provides an opportunity to investigate which optimal external resources can be found or constructed with respect to the cognitive problem in order to enhance human performance.

The notion of scaffolding (Vygotsky 1986; Clark 1998) explains that the learning of a variety of tasks could be made possible, or easier, if suitable external structures and supports are available. Cognitive scaffolding can function by helping to direct attention, and to prompt the right sort of action at the right time, and may be permanently necessary and useful.

A key insight of distributed cognition is that some computational problems can be transformed into intrinsically simpler ones, ones that depend for their resolution on a different mixture of sensory modalities, or ones that can (partly) be solved by means of manipulation. Such transformations can permit multi-modal learning of key relationships. It has been found that transformations can reduce error rates, and increase task effectiveness in various other ways (Kirsh & Maglio 1995).

It has also been found that different ways of encoding or representing the very same information can be more or less helpful for cognition. Cognitively helpful representations make the right sorts of information salient, do not clutter the visual field with irrelevant distractions, help make explicit the structure of what they represent, and can allow what would otherwise be demanding problems of judgement to be replaced by simpler acts of visual inspection (Tufte 2001).

In addition to these motivations from the general field of distributed cognition, the design of our programme is partly motivated by reference to the literature on 'deliberate practice' (see Ericsson & Charness 1994; Ericsson, Krampe, & Tesche-Römer 1993; Ericsson & Lehmann 1996; and the papers in Ericsson [ed.] 1996). 'Deliberate practice' is, it seems, a pervasive feature of elite performers in a wide range of domains including the arts, sciences and various sports and games. Key features of deliberate practice include that it relies on improvement through feedback, that it involves activities focused on single components of performance, rather than the

whole, and that it involves exercises of increasing difficulty. For more remarks on deliberate practice see the section on the UKZN initiative below.

THE *REASON!ABLE* SOFTWARE

The *Reason!Able* software (see Figure 1) is a structured environment for storing and manipulating the components of *arguments*. It is, in different and complementary ways, a form of scaffolding, a transformer of some reasoning operations into manipulations, and a performance enhancing representational system.
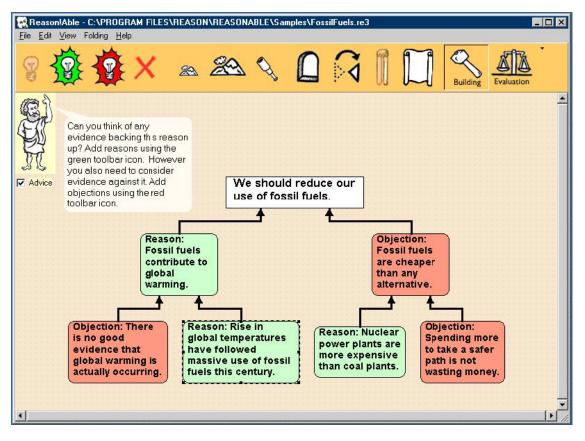

Figure 1: *Reason!Able* screenshot (from http://www.goreason.com)

It is important to recognise and remember that the software is not itself 'intelligent'. It does not 'understand' the arguments built with it, and cannot itself judge their quality. Consequently it is just as possible to fill *Reason!Able* with useless nonsense as it is to do so with a word processor. This is not a limitation – it is crucial to the distinctive value of the system, a point returned to in the conclusion below. Along the lines suggested above, when used properly *Reason!Able* enables users to represent arguments in a way that facilitates rather than impedes understanding and evaluation, and it provides scaffolding that enables more effective criticism and evaluation on the part of the user, crucially including evaluation of her own efforts.

Users of the software build and manipulate representations of arguments, and evaluate the

arguments that have been built. The elements of arguments are *statements* – declarative sentences that can be true or false. Some statements are reasons for (or objections to) others. A well-formed argument has reasons that work together to support the conclusion, just as a proper objection consists of a set of statements that work together to *undermine* a conclusion.

The hierarchical structure of the argument maps built in *Reason!Able* makes the relations between the components (claims that support or undermine this or that claim) visible and explicit through their spatial arrangement, and the lines linking the components. This provides a form of scaffolding. It enables diagnostic questions, such as whether a set of reasons does indeed work together, to be appropriately directed, and (with the support of a few simple rules) makes answering the questions themselves simpler. The system of colour-coding and labelling reasons (green) and objections (red) makes the purported function of particular claims more salient. The fact that parts of arguments can be pulled off the current map, moved to different locations, and transformed (for example, from reason to objection) allows decisions about how to improve the argument map to be carried out by direct manipulation (by means such as dragging and dropping).

Learning to use Reason!Able competently and effectively is not a trivial task. While the environment itself is simple and easy to master, and the key rules regarding how to build and work with good arguments are easy to remember, learning to apply the rules systematically and rigorously takes sustained, structured and appropriately supported practice. It is for this reason that our curriculum development took account of the literature on deliberate practice.

Although learning to use the software effectively is difficult, it is also an excellent way of learning not to make, and of getting into the habit of not making, a range of common errors of reasoning, including those listed above in the discussion of 'the problem'. How is it that correct use of Reason!Able can help with these? The key, in the first instance, is a simple set of rules for correct argument construction. These rules include the following:

Rule 1: *Only* one simple statement per box;

Rule 2: *At least* two co-premises per (sub)-argument

Rule 3: *Absolutely no* 'danglers' (either vertical or horizontal).

Following the first rule ensures that each box in an argument map contains only one statement. Students can fairly easily be taught how to check an argument map to make sure that each box contains a statement (ask "can this sentence be true or false?") and also whether it contains a simple statement or a complex one (ask "can this sentence be separated into two, without adding anything?"). Complex statements (for example, including 'because' claims) are instances of reasoning that has not been made fully explicit, and a sign that work needs to be done in order to make it explicit. Since one of the chief symptoms of insufficiently critical reasoning is failure to make arguments explicit, this simple diagnostic tool has tremendous value.

To understand the second and third rules, an additional illustration will be useful. Figure 2 is an argument map representing a simple and famous syllogism:
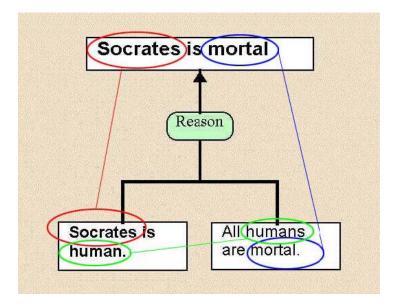
Figure 2: A famous syllogism illustrating the no 'dangler' rule.

Following the second rule ensures that single claims are not allowed to count as arguments, and to prevent repeating a conclusion to be allowed to count as giving a reason. To understand more clearly why this makes a difference, consider the third rule.

A 'dangler' is a part of a statement that only appears once in a sub-argument. If it only appears in the conclusion, then that part of the conclusion is unsupported by its reasons. (These are 'vertical danglers'.) If it only appears in *one* of the reasons, then it is either irrelevant to the conclusion, or it fails to 'work with' the other reasons to support the conclusion. (These are 'horizontal danglers'.)

In the case of the argument in figure 2, there are no 'vertical' danglers, and the parts of the two co-premises that are not connections with the conclusion are connected to each other, so there are no horizontal danglers either. The visible structure of the argument map helps users work out where to look to determine whether the rules are being violated, which is to say that it functions as a form of scaffolding that helps direct attention appropriately.

Repeated application of the second and third rule makes a significant difference to how easily an accurate and/or effective set of relationships between the parts of an argument can be found, as well as helping to direct the process of working out what the components of the arguments are at all. (If an argument made by someone else and recorded in a text cannot be made fully explicit, or does not allow an effective set of relationships between its parts to be constructed, then the failings of that argument are probably being made explicit.) A good argument map enables objections to be accurately directed as well: since the 'no dangler' rule applies in the case of objections too, it can be invoked to help work out what part of an argument is undermined by an objection, and to help make the objection fully explicit.

It seems as though the system works. In Australia, where systematic assessments have been conducted over several years, students taking a single-semester critical reasoning course there show measured gains in reasoning skills over a single 12-week semester of up to *double* that associated with a three-year undergraduate education (van Gelder 2001).

It is less clear what exactly it is that works. Also, the courses at UKZN have yet to be assessed by means of systematic pre and post testing using a standard independent reasoning skills instrument (as the first such testing is taking place in 2005). Some of our students are certainly of the view that something is working, though. In a reflective essay completed at the end of the course (in 2003), one student wrote as follows:

> *As a law student and a future lawyer, it is imperative [for me] to be able to anticipate both sides of an argument. […] Studying philosophy, […] particularly the Reason!Able method of argument mapping, has enabled me to do this more effectively.*

The software also allows initially sketchy representations of an argument to be refined through an iterated process. It is possible to enter a conclusion, and follow it with a series of partial reasons, then once they are in the system, use the rules to determine a better arrangement, and to make the sub-arguments properly explicit.

This system does not, and cannot, make it possible to work with or evaluate arguments without having to make *judgements* at various stages. Nor is that its purpose. Rather, it is intended to provide structured support for the process of determining what an argument is in any given case, and how it is supposed to work, so that judgements are made in a more focused and effective manner. It is possible to gain these advantages with respect to both the arguments of others, and with one's own arguments.

One possible concern that the above account could raise would be that users could form a dependency on the software, so that they could only reason fully effectively when actually using it. The empirical results reported above, in which the assessment instrument was a standard reasoning test completed with pen and paper, and in the absence of the software, suggest that this worry is misplaced. As noted above, some cognitive scaffolding is genuinely temporary – it helps those who lean on it learn the right sorts of relationships, so that their ongoing effectiveness without the scaffolding is better than it would have been otherwise.

THE UKZN INITIATIVE

In our two UKZN courses (both semester-long, one at first-year level and one at second-year level) *Reason!Able* has been integrated into the teaching process of courses that each has its own distinctive content. The first-year course is a general introduction to philosophy, ranging over a wide range of strikingly different topics and historical periods. Several of the outcomes of the course relate to the reasoning skills of students, and include the goals of greater appreciation of the variety of forms of argument, and of the ways in which different forms are appropriate for different problems. The second-year course is on cognitive science, and most of the course is given over to the study of distributed cognition in a variety of domains.

The integration of the course content with *Reason!Able* involves, in the case of both courses, the following components:

- First, all students take a weekly tutorial in a computer room, where a facilitator[2] guides the group through a series of specific exercises and problems, usually involving prepared argument maps and tasks. All of the tasks relate to arguments that are part of the course content, although the first of the 13 tutorials was given over to an orientation to the software itself. Most tutorials involved a mixture of small group and individual work.[3]

- Second, in lectures argument maps built with *Reason!Able* are used as visual aids. This involves talking through the maps, problems faced in constructing them, and commenting on important or interesting features of particular arguments as represented.[4] Over the duration of

      the course a library of argument maps is built up on the lecturer's web page, available for students to download and study, attempt to improve, and/or to extend.

- Third, students are encouraged to use printed argument maps as a point of departure during additional consultations with tutors or the course lecturer.

At the outset of each course in 2003 and subsequent years, time was spent in lectures explaining why *Reason!Able* is designed the way it is, how and why it is supposed to work, and why, although difficult, the tutorial exercises are worth taking seriously. It was make clear that what was being demanded was not easy, and an ongoing effort was made to provide motivation and encouragement. (See the remarks on evaluation below: we may have fallen short in this area.) Since the aim of the course was not primarily to produce students who were good at *Reason!Able* as an 'end in itself', but students who were better at reasoning, none of the instruments of student evaluation (assignments, essays and the final examination) depended on constructing and working with argument maps. Instead the clarity and quality of written argumentation was assessed in line with practices from the preceding year, when the same content[5] had been taught without computer-supported tutorials. That said, students were encouraged but not required to make argument maps during the planning stage of written assignments and invited to submit any such maps along with the assignments for comment and feedback.[6]

A key motivation for the design of the tutorial exercises was, as noted above, to facilitate and guide what Ericsson and his co-researchers call 'deliberate practice'. Key features of deliberate practice include:

- It is designed to provide maximum opportunities for improvement through feedback. (This feedback typically requires an independent specialist source of critical feedback.)

- It involves activities focused on single components of performance, rather than the whole.

- It involves exercises of increasing difficulty, ideally with more advanced stages being attempted only when prior stages have been performed at a pre-determined acceptable level of effectiveness.

- It involves repetition and requires conscious endorsement of the goal to improve.

A striking feature of the research on elite performance across the range of domains studied is that top-level performance (including that of child prodigies) seems to be reached after *ten* years of deliberate practice (Hayes 1981; Ericsson & Lehmann 1996). This suggests that the approach being attempted at UKZN should, if found to be successful, be extended further into the curriculum, rather than thought of as a 'one-off' remedy that takes full effect in one semester.

The tutorial exercises in use at UKZN were specially designed to lead students through a sequence of increasingly demanding tasks involving the application of the rules of reasoning described above. At the simplest extreme they involve checking argument maps for violations of the 'one statement per box' rule, and at their most demanding substantial reorganisation and addition is required in order to repair the argument map or otherwise complete the task. Most of the maps used in tutorials were constructed in a three-stage process. First, accurate maps of arguments covered in the course were built up and refined over an extended period. Second, the maps were deliberately modified in various ways, including deleting premises, telescoping separate claims into one, and moving sub-arguments out of their proper places. Finally, a specific set of written instructions (along with clues and suggestions as to how to proceed) were prepared to accompany the task.

We have a growing library of tutorial exercises, many of them developed following a period of intensive workshopping and experimentation. Materials development is an ongoing priority. Each

tutorial is intended to help deal with a specific reasoning problem, or challenge, or a set of them. Other tutorials in the course were intended to facilitate understanding of arguments where students tend to agree with the conclusion, and hence (sometimes) feel less interest in knowing the exact reasons, or to appreciate how arguments can function as explanations, and hence as a tool for coming to understand complex positions. One of the later topics in the course, taught to humanities students who are often wary of mathematically rich topics, is Einstein's special theory of relativity.

In 2003 and 2004 the primary goal of the curriculum initiative with *Reason!Able* was to develop functioning courses, and, given the demands of meeting that objective, we do not yet have rigorous quantitative data regarding measured improvements in reasoning scores along the lines of the studies reported above. (In 2005 we were required by the demands of an institutional merger to implement a revised curriculum, hence impeding quantitative evaluation further.) Some evaluative comments are nonetheless possible. In what follows I will focus specifically on the first-year course that ran in the second semester of 2003, followed by some remarks on an assessment that is presently underway. First an outline of some of the main demographic features of the class is offered. Second an overview of the sorts of evaluation planned for the future is provided. Finally, preliminary evaluative remarks based on the student course evaluations, and the experience of the tutors, are offered.

The students

In the second semester of 2003 the final enrolment for the course 'Philosophy 1B' was 219 students. Of these students 184 (84%) were in their first year of study, while the remainder of the students were in later years. Forty-four (20%) of the students were drawn from the four-year Bachelor of Business Science programme, and 47 (21%) from the four-year undergraduate Bachelor of Laws programme. Of the remaining students, the vast majority were enrolled for general or programme Bachelor of Arts (46 students, 21%) or Bachelor of Social Science (50 students, 22%) degrees.

Forty per cent of the enrolment consisted of white students, while 56% of the class was female. Seventy-one per cent of the class was under the age of 20 years, and 82% of the group reported English as their home language. Fifty per cent of the class entered university with 40 or more matric points. Success at the course overall correlates strongly with 2002 matric scores – the vast majority of those whom passed the course had 30 matric points or more, and almost every student whose final mark was better than 80% had 40 or more matric points. This correlation is not surprising – a total of 30 or more matric points is known to be a reasonable predictor of success at undergraduate courses at UKZN.

Given that the entrance requirement for the Faculty of Human Sciences is 24 matric points, the high average number of matric points per student in the course requires explanation. The Philosophy programme does not impose its own entrance criteria. Nonetheless the 20% of the group drawn from the Business Science programme have passed entrance requirements calling for 38 matric points including at least a 'B' for Higher Grade mathematics. The 21% drawn from the Bachelor of Laws programme were required to have at least 34 matric points to gain admission to that faculty. Any profile of the class by matric points is limited, furthermore, insofar as some students enrolled via access programmes that attempt to assess prospects for success independent of matric results, or in the case of some students, especially older students, in their absence. Finally, the course has a reputation among students for being relatively challenging (see the extracts from course evaluations below) and this probably plays a role in self-selection on the part of students. We are also well aware that matric points should not be regarded as a context invariant indicator: a relatively low score by a student whose secondary education took place in a very poorly resourced school, for example, could be much more impressive than a

'better' score from a student with access to excellent resources.

The enrolment in 2003, as in every year, included students with various forms of under-preparedness, notably including low critical reasoning skill and low critical reasoning motivation (although sometimes with high general motivation), and in some cases low matric scores. Some under-prepared students achieved striking success – including a few cases in which students with close to, or less than, 24 matric points, and who had attended poorly resourced high schools, achieved final marks in the upper second (68–74%) or first class (75%+) range.

From 2004 this same course is to be offered in both semesters, rather than only in the second. Part of the motivation for this is to allow more students to take the course, but at the same time to work with a smaller number at any given time. It is hoped that the better staff:student ratio will allow more extensive support to be provided to all students, especially those who are under-prepared, but motivated.

Future evaluations

In the first semester of 2005 we began a three-year project to run a series of pre and post tests on classes taking *Reason!Able* tutorials and controls of various sorts. In the first semester of the study we have conducted pre-testing on each of three groups (one at each year of study) in a course using *Reason!Able* tutorials, and will follow these with post-testing and analysis at the end of the semester. This analysis will determine the details of the design of the remainder of the project.

For this study we are using the same instrument used in the Melbourne studies reported above, the California Critical Thinking Skills Test (CCTST). This will enable the results of the UKZN study to be related to those undertaken elsewhere, and will assist in gauging our success. Linking the results of pre- and post-testing with matric points and subjects will contribute in a small way to establishing in more detail in what ways, and to what extent, matric marks overall, and for particular subjects, are predictors of reasoning skill prior to teaching with *Reason!Able,* and also what aspects of matric results, if any, are predictors of capacity to make significant gains over the course of a semester. This study may also be extended to include follow-up work on the varying success of students who completed the course at the variety of programmes of study they go on to follow. Ideally more seriously longitudinal study of a small group of students going through the learning process will be included at some time as well – a noted limitation of pre- and post-testing is that it illuminates so little about what takes place in between.

It would be encouraging to find that the approach adopted at UKZN is producing the same sort of rapid gains as in Melbourne. If it is then follow-up study to determine the extent to which any gains are *retained*, and how they impact on future success, is necessary. Given the research noted above on deliberate practice and expert performance, we also hope to determine the extent to which taking successive courses of this sort (some UKZN second-year courses involve computer-supported tutorials) produces ongoing gains.

Preliminary evaluations

In the absence, for now, of properly independent measures of success of the sorts envisaged and described above, it is still possible to comment on the success of the project in 2003.[7] First, it is my impression (as the person who marked all of the 700-odd assignments and 450-odd examination essays produced in the course) that the level of detail and clarity in written arguments produced by the students improved steadily through the semester, and that the same was true for the quality of verbal arguments in class discussion.

Second, the tutorial facilitators reported that tasks of types that initially had their groups stumped, became easier and easier over time, and that during the second half of the course especially, students in tutorials would regularly complain that the tutorial period was too short, and that they wanted more time to finish what they were doing to their own satisfaction.

Third, as supervisor of the dissertations of two of the facilitators, I can report that the quality of their own written work increased noticeably.

Fourth, the student evaluations of the course included the positive and striking results presented in the Table 1 below. The reason that the results are given in two columns is that approximately 20% of enrolments on the course were first-year students taking Bachelor of Business Science degrees. Timetabling requirements meant that these students took some classes in different venues, and their assessments were collected and processed separately to facilitate a variety of comparisons. The entrance requirements for the Business Science degree are higher than those for all other degrees from which first-year students in the course were drawn, and include higher-grade mathematics. Tutorial attendance records show that the Business Science students attended their tutorials (also held in different venues for timetabling reasons) with considerably greater regularity than other students (on average).

Table 1: Student evaluations of the course

Evaluation question	Positive response – Business Science	Positive response – Humanities, Law, Social Science and others
The course as a whole was "challenging"	89%	90%
The course was "useful in developing my thinking skills"	93%	87%
The course was "useful for (some of) my other subjects"	48%6	87%
The computer-based tutorials were "useful in developing my reasoning skills"	70%	59%
The computer-based tutorials were "useful for understanding the course content"	85%	52%

These considerations suggest that the report of those students who attended the greater proportion of the tutorials may be a better guide to their value – and those students who were most dedicated to attending and participating were mostly convinced that the tutorials were useful. Nonetheless it is not clear that this is a sufficient explanation – rather it raises a new explanatory challenge which is to account for the far more patchy attendance in tutorials of Humanities, Law and Social Science students. One pedagogically lazy explanation suggests that 'good' students simply work harder, and do better. The more challenging, and I think correct, response is to recognise, as empirical research suggests, and as noted above, that overall reasoning effectiveness depends on a combination of skill and motivation. Both need to be encouraged, developed and supported in the teaching process. It is my hope that careful attention to the task of building and consolidating motivation among students will lead to significant gains for all involved. I suspect that some students found the early stages of the tutorial programme somewhat demoralising, and that more could be done to help them perceive the gains that they are making as they go along, and to help them believe that the 'pain is worth

the glory'.

On a personal note, the process of teaching with the software was gratifyingly humbling. It was quite a lot more difficult than I expected to make satisfactory maps of arguments that I'd have thought I knew inside out from years of teaching them. I also discovered that working collaboratively on building maps (which I did with graduate students in the course of making tutorial tasks) was highly rewarding – the maps provide a clear focus for discussion, and an anchor reducing or preventing drift away from the topic.

THE FUTURE

As noted above, 2005 sees a proper quantitative evaluation several of our courses, using the same independent instrument as was used in the Melbourne studies reported above. I hope that it will find that we achieve a comparable gain in measured reasoning skills to that found in Melbourne. In preparation for the 2005 courses a substantial portion of the tutorial tasks have been completely overhauled, including various optional sub-tasks for students who find they are working more quickly than others, and additional extra challenging tasks that can be attempted outside tutorials. (The full set of exercises is to be placed on the web, along with the argument maps that form the basis of each tutorial.)

We are also going to break new ground in the use of *Reason!Able* by integrating it into a number of graduate courses. To replace the common practice of seminars commencing with one person reporting on the content of one of the readings for the seminars, we will require construction of argument maps as preparation for seminars. At the start of each seminar those who have built maps will distribute copies, and then talk everyone through their map, and any difficulties that they ran into along the way. This initial discussion will set up the more free discussion that follows, and hopefully make a significant impact on its quality. The 'ten-year rule' in the empirical study of expert performance reported above suggests that integrating argument mapping into graduate teaching might not be a case of overkill, but could instead continue to provide significant gains. It will be a few years before graduate students in our courses are exposed to deliberate practice in argument mapping in their first year, and hence some time yet before follow-up quantitative assessment of longer-term gains is possible. Such assessment should attempt to measure: (a) the extent of skill retention subsequent to first-year evaluation; and (b) the magnitude of any gain following exposure in the context of graduate seminars.

Following one more year of development on courses in philosophy including considerably more development of online materials, we also hope to be in a position to assist and facilitate the development of curricula in other disciplines using the same general approach. There are precedents elsewhere in fields such as law and nursing, and there is already some interest from other departments and faculties at UKZN.

CONCLUSION

South African universities face their own version of a global problem in the development of critical reasoning skills in their students. While the empirical effectiveness of many standard approaches to teaching critical thinking have modest results, some with proven and remarkable effectiveness are being implemented at UKZN. Preliminary evaluations indicate that the UKZN project is successful, although by no means (yet) as successful as it could be. Future empirical testing will help establish both the degree to which the courses are successful at their intended purpose, and also the extent to which any benefit is retained in the course of further study.

Two points raised above should be returned to in closing. First, it was suggested that a major strength of the approach is that the software used is not itself intelligent. Second, one impediment to the development of critical reasoning was suggested to be a needlessly deferential attitude to perceived authorities. A system based on interaction with a genuinely intelligent computer (were that presently possible, which it is not) would, I suggest, produce just one more authority to which students would defer instead of coming to trust their own capacity to be competent judges. A non-intelligent system that, when properly used, facilitates the clarification, explication and refinement of any argument, including a student's own thinking, is a proper tool for the task of getting more students to take themselves seriously as intellectual actors, competent to challenge any view, rather than mere consumers who attempt to take on the thoughts of others.

Endnotes:

[1] I would like to thank Hanlie Griesel of SAUVCA (the South African Universities Vice Chancellors Association) for encouraging me to present a previous version of this work at a curriculum development workshop organised by SAUVCA in 2004, and Tony Carr and Laura Czerniewicz for encouraging me in turn to participate in the e/merge on line conference also in 2004. I acknowledge the Quality Promotion Unit of the former University of Natal for supporting the initial project, and all of the tutors involved in delivering the initiatives described here. Lynn Slonimsky (Wits), Ian Moll (SAIDE) and an anonymous referee provided useful critical comments on earlier versions of this document. Finally, further thanks to Laura Czerniewicz for comments and advice on the present text.

[2] The tutorial facilitators in 2003 were graduate students, most of whom were trained in cognitive science, and all of whom received specific training in working with *Reason!Able*. Meetings with the tutorial facilitators were held weekly prior to each tutorial.

[3] One limitation of the current UKZN initiative (up to 2005) should be noted here. For budgetary reasons the pilot project is working with a limited-user software licence rather than a full institutional site licence. Consequently the software was installed only in a few computer rooms, and it was not possible legally to provide students with copies they could install and use on home computers.

[4] I regularly subjected my own argument maps from previous lectures to criticism and offered improvements of them. This was partly an attempt to lead by example, and hopefully to discourage any of my own maps from being taken as definitive, instead of regarded critically. It was also an opportunity to work through the application of the three basic rules described above, and illustrate the genuine difficulties that can arise.

[5] The course in 2003 and 2004 was an introduction to philosophy ranging over topics in, *inter alia*, epistemology, ethics, distributive justice, philosophy of science, and political philosophy. As of 2005 the focus has shifted to leave ethics and political philosophy for a companion module.

[6] A fair proportion did so, and some of the resulting argument maps were used to guide the development of further tutorial exercises.

[7] Results from the student evaluations in 2004 were not significantly different from those reported here.

REFERENCES

Clark, A. (1997), *Being There,* MIT Press, Cambridge, Mass.

Clark, A. (1998), Magic words: How language augments human computation. In P. Carruthers & J. Boucher (eds.) *Language and Thought: Interdisciplinary Themes*, Cambridge University Press, Cambridge, pp. 162–183.

Ericsson, K.A. (ed.) (1996), *The road to excellence: The acquisition of expert performance in the arts and sciences, sports, and games*, Erlbaum, Mahweh, NJ.

Ericsson, K.A. & Charness, N. (1994), "Expert performance", *American Psychologist,* 49, pp. 725–747.

Ericsson, K.A., Krampe, R.T., & Tesche-Römer, C. (1993), "The role of deliberate practice in the acquisition of expert performance", *Psychological Review,* 100, pp. 363–406.

Ericsson, K.A. & Lehmann, A.C. (1996), "Expert and exceptional performance: Evidence of maximal adaptation to task constraints", *Annual Review of Psychology*, 47, pp. 273–305.

Hayes, J. (1981), *The Complete Problem Solver*. Philadelphia: The Franklin Institute Press.

Hutchins, E. (1995), *Cognition in the Wild*, MIT Press, Cambridge, Mass.

Kirsh, D. & Maglio, P. (1995), "On distinguishing epistemic from pragmatic actions", *Cognitive Science*, 18, pp. 513–549.

Spurrett, D. (2003), "Why think that cognition is distributed?", *AlterNation*, vol. 10, no. 1, pp. 292–306.

Tufte, E. (2001). *The Visual Display of Quantitative Information,* Graphics Press, Cheshire, Conn.

van Gelder, T.J. (2001), How to improve critical thinking using educational technology. In G. Kennedy, M. Keppell, C. McNaught & T. Petrovic (eds.) *Meeting at the Crossroads*. *Proceedings of the 18th Annual Conference of the Australasian Society for Computers in Learning in Tertiary Education*, Biomedical Multimedia Unit, The University of Melbourne, Melbourne, pp. 539–548.

van Gelder, T. J. (2002), Enhancing Deliberation Through Computer-Supported Argument Vizualization. In P. Kirschner, S. Buckingham Shum & C. Carr (eds.) *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*, Springer-Verlag, London.

Vygotsky, L.S. (1986), *Thought and Language*, MIT Press, Cambridge, Mass.

Original article at: http://ijedict.dec.uwi.edu//viewarticle.php?id=44&layout=html