Enhancing and Augmenting Human Reasoning

Tim van Gelder

Department of Philosophy
University of Melbourne
Parkville VIC 3052
www.arts.unimelb.edu.au/~tgelder

Austhink
Level 50, 120 Collins St
Melbourne VIC 3000
www.austhink.org

Paper presented at *Cognition, Evolution and Rationality: Cognitive Science for the 21st Century*. Oporto, September 2002. To appear in a volume based on that conference edited by Antonio Jose Teiga Zilhao.

Draft 8 September 2003. Feel free to distribute but do not quote or cite. Feedback is welcome.

Contents

Two Pioneers of Computing and Intelligence

If you ask philosophers or cognitive scientists to name the most important pioneer in the general field of computing and intelligence, they would probably pick Alan Turing, the English logician and computer scientist who arguably did more than any other single person to put disciplines such as computer science, cognitive science and artificial intelligence on a solid theoretical footing.

If you asked the same question in Silicon Valley, you might well be told of someone much less well-known in academic circles, even though his impact on the academy has been at least as pervasive. Douglas Engelbart was not a theoretician but an engineer, and his effect has been not on the ideas and debates but on the day-to-day activity of scientists and philosophers, as well as virtually every other “knowledge worker” in the ever-expanding information-based economy.

Engelbart was concerned with the power and productivity of human intellectual activity. Where Turing focused on making computers smart, Engelbart focused on using computers to make us smarter: artificial intelligence (AI) versus intelligence augmentation (IA). His concern, moreover, was with intelligence not in the abstract or in the laboratory, but in real, everyday work situations, dealing with the kind of urgent practical problems upon which the quality of human life directly turns. In a landmark technical report Augmenting Human Intellect: A Conceptual Framework (Engelbart, 1962), he wrote

> By 'augmenting human intellect' we mean increasing the capability of a man to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems. Increased capability in this respect is taken to mean a mixture of the following: more-rapid comprehension, better comprehension, the possibility of gaining a useful degree of comprehension in a situation that previously was too complex, speedier solutions, better solutions, and the possibility of finding solutions to problems that before seemed insoluble. And by "complex situations" we include the professional problems of diplomats, executives, social scientists, life scientists, physical scientists, attorneys, designers—whether the problem situation exists for twenty minutes or twenty years. We do not speak of isolated clever tricks that help in particular situations. We refer to a way of life in an integrated domain where hunches, cut-and-try, intangibles, and the human "feel for a situation" usefully co-exist with powerful concepts, streamlined terminology and notation, sophisticated methods, and high-powered electronic aids.

To meet this challenge, Engelbart and his team invented an integrated cluster of tools and techniques, including the mouse, multiple windows, hypertext linking, integrated text and graphics, and many others. Collectively, these innovations form the foundations of personal computing. They transformed the computer from a symbol manipulator used only by highly trained specialists into a piece of everyday office equipment, used by even the most minimally competent to expedite their ordinary intellectual activities (Bardini, 2000).

About 50 years ago Turing predicted that by around the current time computers would come to be widely accepted as intelligent, as measured by his own famous test (Turing, 1950). This prediction turned out to be well wide of the mark. A decade later, Engelbart predicted that in a similar time frame computers would be widely used to augment human intelligence. This prediction turned out to be right on the money, thanks in no small part to his own contributions.

In Augmenting Human Intellect, he described a system in which statements were entered into the computer using voice recognition, and were then displayed on a screen. With the aid of a keyboard and a "light pen," these statements could be edited, deleted, and spatially grouped; and then links could be drawn to indicate the argumentative structure, i.e., the evidential or logical relationships among the statements. Once this argumentative structure had been displayed, the user had a range of techniques for changing how it was viewed: for example, parts of the structure could be selectively hidden; the user could zoom in or out; some parts could be magnified more the others; and so forth.

In the previous paragraph I described this system in prose, and left it to you (the reader) to imagine the system in action. This was a deliberate if feeble attempt to illustrate the feat involved in conceiving such a system back in 1962. For when Engelbart first described this reasoning-manipulating system, nothing remotely like it existed; indeed, all the essential components (on-screen editing, etc.) which are so familiar to us today, and out of which we might build our own mental conception of the system, had yet to be developed.

Argument Mapping

A core aspect of Engelbart's proposed argumentation support system was the use of node-and-link-type *diagrams* for exhibiting and then manipulating argumentation structure. How did he get this idea? Apparently he was at that time unaware of work done by other people using diagrams to represent the structure of evidence or reasoning.[1] Innovations often result from the insight that an idea which works well in one domain might be carried over and applied in a similar way in quite different domain. For example, a brilliant engineer of an earlier era, James Watt, developed his centrifugal governor for regulating the speed of a steam engine as an adaptation of similar devices already in use for automatic control in windmills (van Gelder, 1995). Presumably Engelbart, as an engineer, was aware of ways in which box and arrow diagrams could be used to assist thinking about complex structures; for example, the use of flowcharts. When he asked how *reasoning* might be represented, these models were already at hand.

Independent work in the diagrammatic representation of complex argumentation goes back decades before Engelbart's proposal. In the early twentieth century, J.H.

[1] D. Engelbart personal communication, 13 Jan 2003.

Wigmore had developed a sophisticated graphical system for analyzing and representing the structure of evidence in legal proceedings (Wigmore, 1913). In this system of nodes and links, a node might represent a piece of evidence such as *J. was author of letter, though it was in a fictitious name*, and a line between that node and others represented its role in a complex structure of argument aimed ultimately at establishing the guilt of the defendant.

To my knowledge Wigmore's legal maps are the earliest occurrence, anywhere in the world, of argument maps, i.e., diagrammatic representations of the structure of natural language ("informal") argumentation. If that is right[2], it is truly remarkable, for the basic idea in argument mapping is very simple. Arguments have structure, often quite complex, and everyone knows that complex structure is generally more easily understood and conveyed in visual or diagrammatic form. That is why, for example, we have street maps rather than verbal descriptions of the layout of cities. This simple principle has been applied in any number of domains far beyond the mapping of spatial layout, yet it was apparently not until the 20th century that the inferential structure of real-world argumentation was graphically depicted. There had of course been many different kinds of diagrams used in various ways in logic and reasoning (Gardner, 1983), with Venn diagrams for syllogistic reasoning perhaps the most well-known example. None of these, however, depicted the inferential relationships between whole propositions of ordinary argumentation, which is the essence of argument mapping; they were mostly concerned with elucidating logical structure of a more fine-grained nature.

Wigmore started something of a tradition within the study of legal reasoning (Anderson & Twining, 1998), though it was at most a minor tributary of legal theory, and almost completely irrelevant to the torrent of actual legal practice. Argument mapping also emerged in various other quarters of the academy, most notably in philosophy with Stephen Toulmin's influential book *The Uses of Argument* (Toulmin, 1958). Interested in argument but dissatisfied with the tools offered by the logical tradition, Toulmin developed a simple diagrammatic template intended to help clarify the nature of everyday reasoning. This template is now widely used for displaying the structure of informal reasoning, especially common in fields such as rhetoric, communication and debating.

Subsequently, argument structure diagrams became commonplace in textbooks and classrooms in the area of instruction in critical thinking or informal logic (Govier, 1988). These argument structure diagrams are usually very simple, and used mainly for pedagogical reasons; they are like training wheels, intended to be thrown away after certain basic principles had been understood and skills (supposedly) acquired.

---

[2] I would be very surprised if there weren't prior examples somewhere, but my own searches and queries have not yet turned up any. If anyone knows of argument maps prior to Wigmore, I would appreciate hearing about it.

They are not used to help sophisticated reasoners think more effectively about complex real-world issues.

The idea that argument mapping might be a tool supporting real thinking, rather than an instrument of academic analysis or an educational stepping-stone, first arose with the work of Robert Horn and his co-workers in the 1990s. His most well-known achievement is *Can Computers Think?* (Horn, 1999), a series of 7 poster-sized charts of the main the lines of argument in some 50 years of academic debate in the philosophy of artificial intelligence – a debate whose starting point was (to close the loop) none other than Alan Turing's claim that computers are capable of thinking and that there will eventually be thinking computers. This massive map is like a large-scale map of a country's road system, in which positions correspond to cities or towns, highways to main lines of argument, and secondary roads to detailed debates. Studying the map, the reader can rapidly and easily identify the main contours of the debate, follow out some particular thread, and rapidly switch between studying the fine detail and seeing the how that detail fits into the larger structure or argument (Holmes, 1999). The maps can help readers overcome one of the greatest challenges involved in research, the process of assimilating the existing literature to the point where one understands how the main arguments relate and where a useful contribution can be made. This process can take many years, and often is never really completed; many researchers have only a foggy sense of the issues and arguments in the field outside their own topic of immediate interest.

Computer Supported Argument Mapping

The *Can Computers Think?* series of maps was a kind of argument mapping megaproject. It involved a team of people working for a number of years researching and then diagramming a very complex and often quite abstruse debate. The maps needed careful design and were eventually sent to a specialist printing service and were then distributed much like any other book. The resulting series of maps have many of the virtues but also many of the problems and limitations of a large printed map of the world:

- They are expensive to produce, requiring time and specialist expertise for background research, map design, printing and distribution
- By their very nature, they cover only one topic. The *Can Computers Think?* maps are a great help if your interest is in whether computers can think, but are not much use at all if it is in whether animals can think or whether President John F. Kennedy was killed by a conspiracy involving the CIA.
- The maps are essentially static objects; you can study them, but not interact with them. Yet interaction is important for understanding and learning.
- They cannot be modified by users. They go out of date, as new research extends the terrain to be mapped and suggests better ways to construe the existing terrain.

For argument mapping to become an everyday form of thinking support, these kinds of challenges would have to be overcome; Horn-style megamaps are valuable but far from the whole story.

This is where the diagrammatic methods of argument mapping connect with Engelbart's vision of the computer as tool for intelligence augmentation. Instead of providing static, prepackaged argument maps, why not provide tools for people to cheaply and easily create their own maps on whatever topic they choose? By the 1990s such tools had become a serious possibility, because the technological infrastructure was already in place. Most "knowledge workers" had or could easily afford to obtain personal computers with monitors, graphics capability, email, and colour printers. Such systems could be used to create argument maps; indeed they were used by Horn and his team. The only missing element was software which would transform the personal computer into a special-purpose argument mapping tool (rather than a generic tool which, with a lot of time and effort, could be used to create argument maps).

There are now a number of such software packages available. Probably the best examples are Reason!Able, Araucaria, and Athena. All were developed in educational or academic contexts, and all are essentially similar in being based around a workspace or canvas upon which "box and arrow" graphs of reasoning structures can be easily drawn and re-drawn.

Enhancing and Augmenting Human Reasoning

The core contention of this chapter is that with the emergence of these new software packages, we are finally starting to see Engelbart's 1962 vision of intelligence augmentation by means of computer-based support for human reasoning being realized. Computer-supported argument mapping is actually starting to help people think more effectively.

This works in two ways, which I call *enhancement* and *augmentation*. First, computer supported argument mapping (CSAM) can *enhance* human reasoning by helping strengthen peoples' intrinsic reasoning skills, i.e., the skills they deploy unconsciously when engaging in reasoning in everyday or professional contexts.

Second, CSAM can *augment* human reasoning by being used "on the job" to help people perform more effectively. In augmentation, CSAM tools are used to extend our intrinsic or unaided capacities, in much the way pole-vaulters can reach much greater heights than high-jumpers through skillful use of a pole. They are not learning aids, but thinking equipment; indeed, in a certain sense, they become part of the mind itself.

Enhancing Human Reasoning: CSAM in Education

Almost everyone, it seems, accepts that one major aim of education is to cultivate thinking skills generally, and in particular the skills of general informal reasoning and argumentation. Unfortunately, aims and outcomes are not always the same. There is incontrovertible evidence that many people emerge from secondary or even tertiary education with general reasoning and argument skills that are under-developed, sometimes woefully so. Perhaps the most substantial body of evidence on this topic is that collected by psychologist Deanna Kuhn and reported in her book *The Skills of Argument* (Kuhn, 1991). Her starting point was the bleak observation that

> Seldom has there been such widespread agreement about a significant social issue as there is reflected in the view that education is failing in its most central mission—to teach students to think. (p.5)

To get a better fix on the problem, she conducted intensive structured interviews with 160 people drawn from a wide range of age groups, occupations, and education levels. She found that while most people are quite able and quite ready to form and express opinions on complex and controversial matters, more than half cannot reliably exhibit basic skills of reasoning and argument in relation to those opinions (and so are not rationally entitled to them). For example, most people will readily hold an opinion as to why some youths stay away from school, but over half cannot provide *any genuine evidence at all* for their position (let alone *good* evidence). When asked something like "What evidence can you provide that your account of why some youths stay away from school is the correct one?" everyone will say something, but in a majority of cases what they say doesn't constitute *evidence*; it might be a restatement of the position, a digression, or an illustration, but not information which would properly induce a rational person to have more confidence in the account.

Anyone with experience teaching undergraduates will recognize the problem Kuhn was diagnosing. A great many students enter university with only feeble understanding, and no mastery, of general reasoning and argument skills. Far too many of those students exit the other end in a similar state of cognitive impairment. Small wonder: overwhelmingly, students get virtually no explicit instruction in these principles and procedures. Their institutions and instructors seem to assume that students either already have the skills, or that they will pick them up by osmosis, imitation and practice, with feedback given haphazardly and in tiny slivers. This works about as well as expecting dog-paddlers to become water polo players without ever showing them how to swim properly.

The most widespread and deliberate approach to confronting this problem is to provide direct instruction in the form of a one-semester undergraduate subject. These subjects, known by names such as *Critical Thinking*, *Informal Logic*, and *Introduction to Reasoning*, are usually provided by Philosophy departments as first-year electives. They aspire to help students to reason better, and are often advertised as having this effect.

However it is far from clear that they succeed. On one hand, there is little serious *positive* evidence of substantial gains in reasoning skills among students taking such subjects. Students' performance on the final test, on its own, cannot be taken as such evidence, since it incorporates whatever competence they brought with them to the subject at the start. To know how much students *improved*, you'd have to measure that initial level and subtract it from their final performance – something that is rarely done.

On the other hand, there is a worrying amount of negative evidence. The preponderance of studies which have attempted in some reasonably rigorous way to quantify the gains attributable to taking such subjects have found that they provide little or no benefit. For example, at the University of Melbourne we used the Watson-Glaser Critical Thinking Appraisal to pre- and post-test students in a conventionally-taught one-semester first-year critical thinking subject. Students performed at essentially the same level in both tests, which was a very disappointing result, considering that they should have improved a modest amount due simply to maturation and being at university. More generally, reviews of studies of attempts to improve critical thinking have tended to pessimism; for example McMillan (McMillan, 1987) reviewed 27 studies and concluded that "the results failed to support the use of specific instructional or course conditions to enhance critical thinking." Pascarella (Pascarella, 1989) found no statistically significant effect of taking logic courses on growth in critical thinking in first-year university students. In fairness, it should be noted that some studies of individual subjects have found gains, and some reviewers have been broadly optimistic (e.g., (Halpern, 2002)). The truth is that empirical studies of growth in reasoning and argument skills at university are the proverbial dog's breakfast, and while it seems clear to me that the overall trend is against any benefit, any reviewer can pick and choose evidence to support their pre-given theoretical perspective (a very uncritical exercise!).[3]

Of course, most instructors *believe* (one hopes!) that their offerings are making some worthwhile difference to skills, and their beliefs are grounded in their direct, informal observations: they can *see* their classes improving under their careful guidance! Further, every instructor can provide plenty of anecdotal evidence; it seems you must be doing something right when you meet a student years later who sincerely attests that taking your class has helped them greatly ever since. But informal observation and anecdotal evidence can also point the other way. My own interest in the topic grew out of my experience teaching Critical Thinking for four years, and giving up in despair. It seemed my students could not have improved much, since they were so

[3] The problem here is the informal literature review – a process whose loose constraints allow reviewers, when presented the sort of heterogeneous body of literature typically found in the social sciences, to select, massage and present evidence so as to support their favoured positions. One way to overcome this problem is to engage in meta-analysis (Hunt, 1999). A proper meta-analysis of studies of critical thinking growth is sorely needed.

bad at the end it was scarcely credible that they were much worse at the start. Similar gloomy assessments have been voiced by veterans in the field. After 30 years trying to teach Introductory Critical Thinking, Doug Walton said

> I wish I could say that I had a method or technique that has proved successful. But I do not, and from what I can see, especially by looking at the abundance of textbooks on critical thinking, I don't think anyone else has solved this problem either. (Walton, 2000)

Further, informal observation and anecdotal evidence are notoriously untrustworthy. Bloodletting survived as a medical practice for thousands of years based in large part on practitioners' direct observation of success with the technique – that is, patients who improved after bloodletting. Of course many patients sickened or died, but this could easily be explained away; after all, the patients were only being treated with bloodletting because they were seriously ill, and even a good technique can't be expected to perform miracles! More generally, informal observation and anecdotal evidence are the foundations of every form of quackery and pseudoscience, from homeopathy and phrenology to Freudian psychotherapy and the selection and remuneration of executives. A form of evidence that has suckered legions of investigators in other fields should not be endorsed by instructors of critical thinking, no matter how sincere and well-meaning they are.

If it is true that conventional instructional techniques are largely or typically ineffective, what might be done to improve the situation? The discussion above, of argument mapping, suggests an obvious possibility. The whole point of argument mapping is to display the structure of reasoning and argument more clearly and explicitly, so that the reader can follow the reasoning more easily and thus think more effectively about the issues. Perhaps, if argument mapping were to form the basis for *instruction* in reasoning, students would find logical structure easier to grasp and logical procedures easier to master.

This idea is hardly original. For decades, many teachers of informal logic have believed that studying and producing diagrams of arguments can help students understand the structure of reasoning, and thereby help improve their reasoning and argument skills. Indeed, at least as far back as Michael Scriven's classic *Reasoning* (Scriven, 1976), it is quite standard for introductory textbooks have a section on argument diagramming.

However, these diagramming activities are generally limited to simple argument structures with at most a handful of nodes, and play only a minor role in the overall pedagogical approach.[4] Why is this? Here are some conjectures:

---

[4] In this tradition, Alec Fisher's textbook *The Logic of Real Arguments* (Fisher, 1988), lies at one extreme; it takes on the challenge of handling "sustained theoretical arguments," and deploys structure diagrams which sometimes have a dozen or more nodes.

- Instructors and textbook authors *assume* that very simple diagrams are sufficient to give students the general idea, and that there is no additional benefit in continuing to use diagrams for more complex structures.
- Instructors and students have not had good tools for producing, manipulating and distributing argument diagrams, and so diagramming has been slow and inconvenient.
- Diagrams of the sort normally used have at least one severe usability problem. In these diagrams, claim is usually represented by a label such as a number, often placed in a circle ("box"). Somewhere outside the diagram, the correspondence between labels and claims is set up. The person reading the diagram must hold these correspondences in their mind; they must, for example, remember that claim 3 is *Smith asserted that Jones was not in the house at the time of the robbery*, while claim 4 is *Jones answered the phone at the house soon after the robbery*. But our minds are severely limited in their capacity to maintain such internal databases. The cognitive burden becomes increasingly unmanageable as the number of nodes increases beyond a small handful. Beyond this point, the diagram either becomes opaque, or the reader must continually refer elsewhere to find out what claim a particular number refers to. Given this mental labour, argument structure diagrams of this sort rapidly become useless as their size increases.

The latter two considerations are practical problems, and can now be largely overcome using the new argument mapping software packages; and the first consideration is only an assumption, maintained largely because in practice it couldn't be tested. So, can instruction based on argument mapping be applied even to quite complex structures substantially boost reasoning skills?

The Reason! Project at the University of Melbourne has been investigating this question. We have produced a novel design for a one-semester subject, in which students undergo a structured training regime consisting of many exercises which almost always involve constructing and manipulating argument maps. These mapping activities are supported by a specially-built software package, Reason!Able. The software makes it possible for students to:

- **Rapidly assemble** diagrams by (a) using simple point and click operations build an argument tree, and (b) typing the claims into the relevant boxes. The software "scaffolds" these construction activities in the sense that students can create *nothing but* argument structures; these structures are "syntactically" well-formed even if their content is confused or incoherent.
- **Modify** argument diagrams by deleting nodes, ripping nodes off the tree, drag-and-dropping nodes or even whole branches into different positions, etc.

- **View** argument diagrams in various ways. The software allows students to pan (scroll around), enlarge and reduce, zoom in on particular parts, simultaneously view a part and the whole, and to rotate the entire structure. These options allow students to rapidly reposition their viewpoint so as to maximize comprehension of the complex structure.
- **Evaluate** arguments, and represent their evaluations by superimposing on the argument structure a layer of colours standing for relevant logical qualities. For example, the likelihood that a particular premise is *true* is represented by shades of blue or gray; the strength of a whole reason[5] is represented by a shade of green. In this way the argument diagram displays both structural and evaluative information in the one place; it presents relevant information with a density and immediacy just not possible in prose.
- **Distribute** argument diagrams in standard ways – e.g., copying diagram images into documents or presentations, emailing files, placing files on a websites, and printing out as many copies as desired, perhaps in full colour.

There are two main kinds of exercises in the Reason! approach. In a *critical evaluation* exercise, the student takes an argument as presented by somebody else, identifies its structure and then evaluates it. The student then compares the resulting diagram with a "model answer," i.e., a diagram produced by the instructor; this provides a kind of feedback on the success of her attempt. In a *production* exercise, the student develops an argument of her own; this involves not only working out the structure of the argument, but evaluating that argument in order to determine that it is in fact solid.

Does the Reason! method work? To answer this we have conducted a series of studies in which students are pre- and post-tested using an objective test, the *California Critical Thinking Skills Test*. In study after study, the average gain is around 4 points or about 20% of the pre-test score. In other words, students consistently perform about 20% better on this particular test at the end of the one-semester of training based on argument mapping.

A 20% gain may not sound like much, but in generic cognitive skill acquisition, this is substantial. To show this, we have to first convert the results to the *lingua franca* of empirical studies of the impact or effectiveness of some intervention or technique, known as the *effect size*. The effect size is just the *difference* expressed as a proportion of the natural variation in the population, or what is technically referred to as the *standard deviation*. In the case of studies of critical thinking instruction, the

---

[5] In the Reason! theory of argument structure, a reason is a complex object which always has multiple premises; these are what informal logicians usually refer to as linked or convergent premises.

effect size is the average gain divided by the standard deviation of their performance on the test.[6]

When we make this conversion, we find that students learning with the Reason! method consistently show gains with an effect size of around 0.8. How big is this? Here are some comparisons:

- Cohen, the statistician who developed the concept of effect size, suggested that we might informally think of an effect size of 0.3 as small, 0.5 as medium, and 0.8 or above as large (Cohen, 1988).
- Take a group of students enrolled in a typical first-year non-critical thinking subject (e.g., a political science subject). Most of these students would be getting no direct critical thinking subject. Averaging over all available studies, we find that they gain around 0.3SD over one semester. (This increase plateaus after first year; in other words, after first year, students improve their skills only very slowly.) Alarmingly, this is about the same as the average gain over many studies of students enrolled in a subject where there is *some* critical thinking instruction, which is itself about the same as the gain for students enrolled in a full, dedicated, one-semester critical thinking subject. In other words, standard critical thinking instruction *adds no value* to the overall experience of being an undergraduate. The Reason! approach is therefore much more beneficial than such instruction.
- Based on other studies, students normally gain around 0.8SD over an entire undergraduate education. In other words, the Reason! method can compress the expected benefit of university education, in terms of critical thinking skill development, into a twelve week period.
- In IQ measurement, a standard deviation is 15 points. Thus, if students were to increase their IQ by the same magnitude, they would be gaining one point per week for twelve weeks.

In short, students using the Reason! method do make strong gains. Importantly, these gains seem to hold up over time. When we re-test the students a year down the road, their skills are at the post-test level; they haven't "regressed" back to their starting levels. Instruction based on argument mapping seems to have produced a stable, permanent elevation in their critical thinking ability.

Many reports of a more anecdotal nature support the general view that computer-supported argument mapping really does lead to enhanced general critical thinking skills (Twardy, forthcoming).

---

[6] When comparing different studies, it is more correct to use the "population" standard deviation for the test (the one found in the manual produced by the developers of the test) rather than the standard deviation of the particular group of students in the study.

Augmenting Human Reasoning: CSAM in Professional Practice

When designing and building the Reason!Able software, we assumed that it was just an educational tool. Like training wheels on a bicycle, it would be useful in an early phase of learning, but would eventually become useless or even a hindrance, and so would be left behind.

However, something quite unexpected happened when students used the software to practice their critical thinking skills on increasingly elaborate arguments. It became apparent that the maps produced in the software, and the interactions it supported, made the arguments much easier to understand. In other words, the software-supported argument mapping seemed to be extending the capacity of students handle the sort of complexity found in real-world deliberation.

This suggested that computer-supported argument mapping might be useful outside the educational context, whenever somebody tries to think their way through some complex set of arguments. In particular, it might be useful for people working in professions, government, or business who must regularly engage in complex reasoning activities. Used this way, CSAM would not be *enhancing* an individual's basic reasoning skills. Rather, it would be *augmenting* their basic capacities, leveraging the power of their biological thinking apparatus. CSAM would not be an educational stepping-stone, but rather an everyday tool or piece of equipment, extending thinking capacities just as construction equipment extends our building capacities.

To understand how this might work, consider just one way in which CSAM can be used to augment intelligence – in this case, the collective intelligence of a group of people attempting to achieve rational resolution on some contentious issue through argumentation. It is quite common in organizations for a team or committee to get together to argue things out, i.e., to engage in a kind of collective deliberative process in which opinions and arguments are expressed, objections raised, and so forth. This often takes place in a meeting room, with group members arranged around a large table, perhaps with the most senior person at the head; somebody opens the debate and then it is an argumentative free-for-all; at the end of the meeting, if all goes well, there is some consensus over the core issues.

An alternative to this standard practice is to use real-time CSAM as the framework within which to represent and conduct the disputation. In this approach, the participants still gather in a room, but attention is focused on a screen or wall, onto which is projected an image of an argument map representing the current state of the debate. The map evolves as participant make their argumentative "moves". Collective deliberation mediated by real-time CSAM has a number of advantages over its traditional counterpart:

- **Expanded Grasp**. Individual participants are able to comprehend more of the relevant arguments. Instead of having to construct and maintain elaborate representations "in the mind" – a very laborious and error-prone activity – participants can rely on the argument map as an up-to-date representation of the state of the debate, and scan the map as necessary in order to maintain understanding of both the overall structure of the argument and any particular part of it. Indeed, the map could be said to have become their mental representation of the debate, a representation which merely happens to be located outside the head.
- **Common Mental Representations**. Normally, each participant can remember or hold in mind only a part of overall debate, and each participant retains a somewhat *different* part. This causes tremendous inefficiencies in collective deliberation. When the deliberation is supported by argument mapping, all participants are attending to the one argument map, and so they have a *common* mental representation of the arguments. At that point, their minds could be said to overlap (they are "of the same mind").
- **Targeted contributions**. In ordinary face-to-face argumentation, contributions to the debate fly in all directions; often it is hard to see exactly where somebody's brilliant argument actually fits, and what difference it makes. Using argument mapping, however, contributions can be required to be targeted at a specific place on the map, and to make some specific contribution at that place (e.g., providing evidence against a particular proposition).
- **Depersonalized debate**. An endemic problem in collective deliberation is that positions and arguments tend to be associated with particular people, introducing a range of emotional and social considerations which interfere with cool-headed rational evaluation and constructive engagement in the debate. CSAM defuses this problem by having participants attend to the map and the arguments represented upon it rather than the people making the arguments. It thus "de-personalizes" the debate, making the whole exercise more productive.
- **Organizational memory**. Real-time CSAM generates a large map representing all the arguments, a map which then forms part of the organization's collective "memory" of the debate. If some rational consensus is reached, the grounds for that consensus can be accessed more easily by referring to the archived map than by using unaided recollection or other traces such as minutes, summaries or reports.

Given these apparent advantages, does real-time CSAM in fact improve the overall quality of collective deliberation? It is certainly plausible that it does, but to my knowledge, there is no hard evidence on this topic. It simply hasn't been investigated properly (though see various chapters in (Kirschner, Buckingham Shum, & Carr, 2002)). There is plenty of informal or anecdotal evidence that CSAM-supported deliberation is in fact superior, but such evidence should be taken as at most suggestive.

Real-time CSAM supporting group deliberation is only one way in which CSAM can be used in the workplace to augment human capacities. Others include:

- **Case development**. Individuals, teams or organizations often produce positions (policy statements, recommendations, conclusions) which must be defended by elaborate, convincing argumentation, i.e., there must be a solid *case* for the position. CSAM can be used to develop the case with greater clarity and efficiency than is usually possible, leading to stronger arguments and better-structured documents presenting those arguments.
- **Communicating Complex Arguments**. Argument maps are dramatically more efficient than ordinary prose for communicating the structure of complex arguments. Thus, argument maps can be used to increase the efficiency of communication. For example, currently in law firms around the world, partners frequently commission junior lawyers to draft letters of advice for clients. These letters present arguments, often involving quite complex reasoning, in traditional prose format. The partner must identify and critique the reasoning before the advice is released to the client. This process is slow, difficult and unreliable. An argument map presents the same reasoning but in an easily assimilable, unambiguous form, allowing much more rapid comprehension, evaluation and feedback to the junior lawyer. The time in resources involved in producing the final written advice for the client can thus be substantially reduced.
- **Critical Review**. When an organisation releases a particularly controversial report, it likes to be sure that it has rock-solid arguments to back up its position. Conversely, when critiquing a report whose message is unwelcome, it is useful to be able to identify the arguments so as to be able to target criticisms with greatest effect. In both cases, there is a process of critical review, in which the argument presented in prose is assessed for quality. Argument mapping is a way of exposing the arguments for critical scrutiny, so that the review process can have maximum impact.
- **Design Rationale.** When a team designs a complex artefact, such as a new car or piece of software, many design decisions must be made along the way. These decisions are usually made by a process of collective deliberation. In other words, the team members get together and consider the arguments for and against particular choices. Design rationale is the process of supporting design decisions with clear, strong and persuasive arguments, and recording those decisions for later review. CSAM techniques can support design rationale by providing simple, transparent displays of the reasoning behind a given decision (Buckingham Shum, 1996).

Augmenting reasoning in professional contexts through the use of CSAM is quite new, and has been used in only a few organizations. This can be explained, in part, by the limitations of the available tools. The current CSAM technologies are just crude first steps in the direction of the sort of thinking support tools which will, eventually, transform organizational practices.

Embodied Reason

Assuming that CSAM *does* effectively enhance and augment human reasoning, why does it work? How does it help? What is it about CSAM which makes the difference?

The answer, I think, is that CSAM makes reasoning more "embodied" than do traditional practices and technologies. To see this point, note first that evolution did not bequeath homo sapiens with special, dedicated apparatus for high-level thinking. Rather, we evolved cognitive machinery for other purposes, machinery which just happened to be recruitable for increasingly sophisticated thinking activities, including reasoning and argument. CSAM works because it taps into these available resources more directly than traditional, prose-based ways of engaging in reasoning. In other words, it exploits our embodiment more effectively than traditional modes of argumentative expression.

What, after all, is our brain for? Its primary task, after maintaining basic bodily functions, is coordinating movements in relation to the opportunities and threats in our physical environments. The cognitive machinery evolution has provided is largely devoted to such coordination. We must be able to tell what is around us and rapidly respond to what we find. Hence we have exceedingly fast and powerful hard-wired capacities to distinguish visual features such as line, colour, shape, and position in space, and to recognize the complex, shifting patterns these features can generate. We also have extensive circuits whose primary role is to control and direct the muscular activity whose aggregate expression is the movement of our limbs, the relocation of our whole bodies, and the manipulation of objects around us. These actions play out in sequences whose order and purpose amount to intricate dances with our changing environments.

Higher cognition is possible not because God, evolution or any other benign creator granted us immaterial minds or dedicated neurobiological apparatus. Rather, it is possible just insofar as we are able to redeploy the basic equipment whose primary role is the more mundane business of maintaining and moving our bodies in relation to our surrounds. We can reason not because we have special ratiocination modules but because our sensorimotor control loops can be requisitioned and ordered to operate in fields of grammatically structured and evidentially interrelated objects. This is certainly a remarkable development, but it is evidence not of custom-built mental tools for logic but rather of small transition which happened to nudge our brain's existing equipment over a threshold and into a new range of capacities.

Computer-supported argument mapping is effective because it allows us to redeploy this equipment to the reasoning task more effectively. When identifying or comprehending argument structures, we can take advantage of our ability to recognize line, colour, shape and position in space simply because argument maps, unlike prose, utilize line, colour, shape and position in space to convey information about argument structure. When constructing or modifying argument structures, we can take advantage of our ability to move our limbs and objects around us simply because computer-supported argument mapping supports virtual versions of these basic physical activities. In short, reasoning supported by CSAM is far more like playing with blocks in the kindergarten, or playing with rocks in the vegetable garden, or chasing prey on the savannah, than is reasoning supported by spoken or written prose. The abstraction and complexity of evidential structure is translated, as much as possible, into concrete forms mirroring the particular set of primitive capacities which constitute homo sapiens' evolutionary endowment. Embodied interaction with argument structures gives us new powers to generate, comprehend and communicate complex argumentation. We become more "at home" with complex reasoning because we are already deeply "at home" with our bodies and our surrounds.

The Future of CSAM

CSAM is less than two decades old. Software which supports argument mapping in practice and not just in theory – software that is genuinely usable by people other than the developers – is less than a half-dozen years old. The field is still in its infancy. Inevitably, and rapidly, current technologies will be superseded in every dimension. It is worth speculating briefly about where this technological evolution will take us.

An interesting glimpse of the future was provided in the recent movie *Minority Report*. In that movie, the "pre-cognitions" (visions of future events) of a special group of seers were displayed on a very high resolution, semi-circular, room-sized screen. Large quantities of image and text were manipulated by a character with no physical connection to the screen or computer other than "light gloves" on his hands. Using all the degrees of freedom afforded by two arms and hands, the character could direct the information flow by elaborate bodily performances whose closest analogy would be the motions of a conductor directing a full orchestra in the climax of a late romantic symphony. The information processing in *Minority Report*, referred to as "scrubbing the image," was not argument mapping, though it did involve reaching conclusions on the basis of slivers and patterns of evidence.

The only problem with the *Minority Report* scenario is that it purports to describe the state-of-the-art in information processing technology in the year 2054. This is a very timid vision. In fact, many of the components are already available in "off the shelf" versions. Setups of the kind imagined in the movie will be with us, even common, within a decade.

By 2054 – perhaps well before - high-resolution “displays” will have been miniaturized and implanted inside our skulls, feeding their signals not via light into the retina but via electrical input into our neural circuitry. Our mental control of these displays will not be mediated by motions of the arms and cumbersome technologies such as the mouse; rather, it will be direct “thought” control as the display reacts to the cloud of electromagnetic signals generated by our thinking processes. Artificial intelligence (assuming this is available in only twice the period Turing originally envisaged) will automatically and invisibly aid us in generating and manipulated the displayed information. And of all this, we will be completely unconscious, just as when thinking through a philosophical problem, we are completely unconscious of the details of the operation of our brains. The CSAM technology will have become a permanent prosthetic extension to our biological capacities, a seamless “mindware-upgrade” (Clark, 2003) no more remarkable than, today, a pair of spectacles is a vision upgrade. CSAM will be just one among many respects in which computer technologies will invisibly augment human cognition by providing resources and capacities which complement the strengths and weaknesses of our biological equipment.[7]

Argument Mapping, Rationality and Human Nature

In his weighty tome *Making It Explicit* (Brandom, 1994), Robert Brandom selects rationality as the most profound principle of demarcation between “us” and other things:

> We…[are] the ones who say ‘we’…Saying ‘we’ in this sense is placing ourselves and each other in the space of reasons, by giving and asking for reasons for our attitudes and performances... [it is] to identify ourselves as rational – as the ones who live and move and have our being in the space of reasons… (pp.4-5)

We are the ones who, as he puts it elsewhere, *play the game of giving and asking for reasons*. It is playing this game which makes us what we are; and since we are the ones who created the game – the designers, players, and referees – we are self-constituting beings:

---

[7] Anyone who doubts the general plausibility of the speculations should stay tuned to The Harrow Technology Report, http://www.TheHarrowGroup.com. My conjectures will surely be wrong in detail, but we can be equally sure that something just as dramatic will be true.

> this expressive account of language, mind, and logic… is an account of the sort of thing that constitutes itself as an expressive being – as a creature who makes explicit, and who makes itself explicit. We are sapients: rational, expressive – that is, discursive – beings. But we are more than rational expressive beings. We are also logical, self-expressive beings. We not only make it explicit, we make ourselves explicit as making it explicit. (p.650)

It almost goes without saying that this “having our being in the space of reasons” is physically mediated. Playing the game requires us to *do* things, and these doings are physical processes. For example, to be a rational, discursive agent, you must be able to *say* things – or more generally, to make *assertions*. We humans most commonly, and primordially, make assertions through bodily motions such as exhalations and flexing of vocal cords. Without *some* such physical activity, it would be impossible to participate in any way.

Since playing the game is physically mediated, the nature of the mediation can be varied and possibly improved. We can introduce new tools or *technologies* which alter and perhaps extend our capacity for rational engagement. Perhaps the most obvious and momentous example is the development of writing. Writing allowed us to produce stable representations of moves in the game. At any given time writing is supported by a particular set of technologies (e.g., quills and parchment), and over time these technologies have improved (ball-point pens, now word-processors and the world wide web). The deployment of ever-more sophisticated supporting technologies allows us to play the game more effectively, to make moves that are more informed, more responsive, more nuanced and more widely communicated. We can, quite simply, perform at a higher level than was possible when our physical means were more limited. Of course, this outcome applies to us collectively; it does not imply that any one of us is a more powerful player than was, say, Socrates or Aquinas. As a group, however, our rational self-expression is vastly richer than that of more primitive societies, just as contemporary Rugby Union is a vastly more elaborate and skilful exercise than the anarchic scrimmages from which it emerged.

CSAM is a late development in the history of humanity’s attempts to improve the game of giving and asking for reasons by engineering new ways to play it more effectively (Monk, 2001). It some ways it is an incremental advance, bundling recent technologies and established methods into a new package. Yet incremental advances - such as the adaptation of the wine press into a printing press – can sometimes induce dramatic transformations in the way the manner and level at which intellectual performances are conducted. It is too early to be sure, but this may turn out to be the case with computer-supported argument mapping. For the first time, it seems, we can effectively escape the tyranny of prose as the obligatory medium of argumentative expression. Graphical displays of argument structure are patently better suited to our biologically-given cognitive architectures than prose presentations; properly designed software and hardware are giving us the ability to produce, manipulate and distribute these displays as easily, if not more so, than we can craft argumentative prose.

In other words, CSAM is emerging as the most effective means of *making explicit* our reasonings and deliberations. And if, as Brandom argues in his unique and almost impenetrable style, in making explicit we make *ourselves* explicit, CSAM will be a medium through which humans can self-realize more profoundly than ever before. Since playing the game of giving and asking for reasons is what makes us what we are, playing the game better – or playing a more sophisticated game – is an important stage in the evolution of our nature as sapient beings. Engelbart's vision of computers augmenting human intelligence is, properly understood, a vision of human self-transformation through a bootstrapping process in which our current, technologically augmented intellectual capacities enable us to refashion the spaces and practices within which we ontologically self-constitute. Moreover, his crucial insight was that computer technology will be more profoundly and intimately connected with the process of self-constitution through enhanced rational self-expression than any previous technological forms.

References

Anderson, T., & Twining, W. (1998). *Analysis of Evidence: How to Do Things with Facts Based on Wigmore's Science of Judicial Proof* (Revised ed.): Northwestern University Press.

Bardini, T. (2000). *Bootstrapping: Douglas Engelbart, Coevolution, and the Origins of Personal Computing*. Stanford: Stanford University Press.

Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*: Harvard University Press.

Buckingham Shum, S. (1996). Design Argumentation as Design Rationale. In A. Kent & J. G. Williams (Eds.), *The Encyclopedia of Computer Science and Technology* (Vol. 35, pp. 95-128). New York: Marcel Dekker.

Clark, A. (2003). *Natural Born Cyborgs: Why Minds and Technologies Are Made to Merge*. Oxford: Oxford University Press.

Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale NJ: Lawrence Erlbaum Associates.

Engelbart, D. (1962). *Augmenting Human Intellect: A Conceptual Framework*. Menlo Park, CA: Stanford Research Institute.

Fisher, A. (1988). *The Logic of Real Arguments*. Cambridge: Cambridge University Press.

Gardner, M. (1983). *Logic Machines and Diagrams* (2nd ed.). Chicago: University of Chicago Press.

Govier, T. (1988). *A practical study of argument* (2nd ed.). Belmont, CA: Wadsworth.

Halpern, D. F. (2002). *Thought and Knowledge: An Introduction to Critical Thinking* (4th ed.). Hillsdale, N.J.: Lawrence Erlbaum Associates.

Holmes, R. (1999, 10 July). Beyond Words. *New Scientist,* 32-37.

Horn, R. E. (1999). *Can Computers Think?* : MacroVU, Inc.

Hunt, M. (1999). *How Science Takes Stock: The Story of Meta-Analysis*: Russell Sage Foundation.

Kirschner, P. J., Buckingham Shum, S. J., & Carr, C. S. (Eds.). (2002). *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. London: Springer-Verlag.

Kuhn, D. (1991). *The Skills of Argument*. Cambridge: Cambridge University Press.

McMillan, J. (1987). Enhancing college student's critical thinking: A review of studies. *Research in Higher Education, 26*, 3-29.

Monk, P. (2001, 16 March). Mapping the future of argument*. Australian Financial Review,* pp. 8-9.

Pascarella, E. (1989). The development of critical thinking: Does college make a difference. *Journal of College Student Development, 30*, 19-26.

Scriven, M. (1976). *Reasoning*. New York: McGraw-Hill.

Toulmin, S. (1958). *The Uses of Argument*. Cambridge: Cambridge University Press.

Turing, A. (1950). Computing machinery and intelligence. *Mind, 59*, 433-460.

Twardy, C. (forthcoming). Argument maps improve critical thinking. *Teaching Philosophy*.

van Gelder, T. J. (1995). What might cognition be, if not computation? *Journal of Philosophy, 91*, 345-381.

Walton, D. (2000). Problems and useful techniques: My experiences in teaching courses in argumentation, informal logic and critical thinking. *Informal Logic, 20*, Teaching Supplement #2, TS-35-38.

Wigmore, J. H. (1913). *The principles of judicial proof : as given by logic, psychology, and general experience, and illustrated in judicial trials*. Boston: Little, Brown.