

The Rationale for Rationale™

Tim van Gelder¹

Abstract. Complex reasoning and argumentation are central to legal practice. Software-supported argument mapping may be able to help lawyers reason and argue more effectively. This article describes Rationale™, a generic argument mapping software package, and reviews some evidence that using it can help improve reasoning, i.e., make people smarter. It then explores three different explanations for this potential benefit: usability, complementation, and semi-formality. First, argument mapping software can be more usable for reasoning activities than traditional methods because it can inherit the wisdom gained through decades of research and experience into usability; can exploit a wider range of representational resources; and is designed specifically to support reasoning activities. Second, such software works by complementing the strengths and weaknesses of our natural or inbuilt cognitive capacities. Third, it helps shift reasoning and argumentation into a semi-formal mode, a kind of “sweet spot” between the laxness of everyday reasoning and the straightjacket of formal logic.

1	Introduction	3
2	Rationale Overview	4
3	Making People Smarter.....	11
3.1	Educational Use.....	11
3.2	Professional Use	13
4	Why does it work?.....	16
4.1	Usability	16
4.2	Complementation.....	21
4.3	Semi-formality	26
5	Conclusion - Rationale and Legal Reasoning.....	31

¹ Department of Philosophy, University of Melbourne; and Austhink Software Pty Ltd. Disclosure: the author has a commercial interest in Austhink, which develops and markets the Rationale software.

The power of the unaided mind is highly overrated. Without external aids, memory, thought, and reasoning are all constrained. But human intelligence is highly flexible and adaptive, superb at inventing procedures and objects that overcome its own limits. The real powers come from devising external aids that enhance cognitive abilities. (Norman, 1994)

1 Introduction

The arts of reasoning and argumentation are at the heart of what lawyers do. Lawyers are skilled in these arts, but the complexity of legal issues and the pressures of legal practice call for ever-greater quality and efficiency.

Our efforts over the past decade have been aimed at improving the human capacity to reason. One result of these efforts is Rationale, a software package intended to help people reason more effectively in two ways: by improving fundamental skills, and by supporting the application of those skills to real thinking challenges.

Rationale was not designed for legal reasoning. However there is incipient use in law schools, legal training programs, and "on the job" in law firms. We expect this trend to strengthen as the profession becomes more aware of the potential benefits of such tools, and the tools themselves become better adapted to the needs of legal practice.

This article provides an overview of Rationale; reviews some evidence that using it does in fact help people reason more effectively; explores three explanations as to why it helps, insofar as it does; and concludes with some brief observations relating Rationale to legal practice.

Rationale is one instance of an emerging category of software applications intended to improve general reasoning and argumentation through visualisation or "argument mapping" (Kirschner, Buckingham Shum, & Carr, 2002), a number of which are discussed in other papers in this volume. Although the focus in what follows is on Rationale, the lasting interest, if any, will be in the more general points it makes about argument mapping and how suitably-designed software might enhance human reasoning in general and legal reasoning in particular.

2 Rationale Overview

Rationale is, in essence, a tool for diagramming reasoning on any topic. It supports rapid building, modifying, viewing, and sharing of these diagrams. The term "argument mapping" is loosely used to describe such activities. Thus, Rationale is an argument mapping software package.

Reasoning, no matter what the topic, consists of traversing the web of evidential² relationships among propositions. Hence any diagram of reasoning must somehow represent propositions and relationships among them. The most obvious and common way to do this is with "box and arrow" or "node and link" diagrams. Rationale helps users produce box-and-arrow argument maps.

Rationale is built around a "workspace," an infinite virtual surface on which argument maps can be displayed. Units of information or meaning are created on the workspace and assembled into structures. The Rationale development team uses the quasi-technical term "infons" for these basic units, but for Rationale users they have various more familiar names such as "item," "claim," "position," or "premise," depending on context. By default infons are displayed as rectangular white text boxes with rounded corners, but their visual appearance changes depending on the role they are playing.

Drag-drop operations are used to assemble infons into "thought structures" of various kinds. At this time there are three kinds of thought structures, called "maps" - Grouping, Reasoning, and Analysis.

Grouping Maps

Grouping maps organize infons into hierarchies or trees, displayed in "pyramid" style with the root at the top. Strictly speaking, Grouping maps are not argument maps, but these sorts of hierarchical structures can be effective aides to thinking, as can be seen by the widespread use of pyramid diagrams (Minto, 1995) and mind maps (Buzan & Buzan, 1994). Due to the lack of

² In this paper the word "evidential" and its variations are used with their ordinary meaning rather than with the technical meaning they would be given in many legal contexts.

any constraints other than hierarchy, these maps can be used to help structure thinking for an indefinite variety of purposes and topics.

Reasoning Maps

Reasoning maps are superficially similar to Grouping maps in having an hierarchical structure. However their sole purpose is to display, in the most simple and accessible way, the body of evidence bearing upon some proposition. Thus, the links in a Reasoning map are evidential, and the infons become Positions, Reasons (positive or supporting evidence) or Objections (negative or opposing evidence).

An infon's role in reasoning is signified by various visual attributes. Thus when functioning as a Reason, an infon is green and has rounded corners on both the box and the link to its parent; when functioning as an Objection infon it is red, has sharp corners, and a flat connection between link and box. An interesting variation is the rebuttal - an objection to an objection - which is visually very similar to a standard Objection except that it is colored orange.

Reasoning maps also introduce "Basis boxes" representing various kinds of non-propositional end-points of lines of evidence. In argumentation, the main maneuver is to support (or oppose) claims with further claims. However this cannot go on forever, for both theoretical and practical reasons. In practice argumentation generally and provisionally "bottoms out" in non-propositional material of one kind or another, material which potentially provides evidential (though not, strictly, logical) support. Basis nodes in Rationale provide a way to indicate the presence of such non-propositional grounds.

Reasoning maps are deliberately very simple. They are intended to be useful for naive users, or for more sophisticated users who want to do some quick mapping or who just want a compact distillation of complex reasoning. Representing the true structure of reasoning requires a more sophisticated kind of diagram.

Analysis Maps

A key feature of reasoning, when analyzed carefully, is the notion of linked premises - i.e., claims which "work together" to support or oppose some other claim. Logicians and argumentation theorists have found this notion hard to pin down, in the sense that it is difficult to specify in a general and rigorous manner when claims should be regarded as linked premises

of a single argument as opposed to premises of distinct arguments (see, e.g., (Snoeck Henkemans, 2000) for a review of various attempts). However, it is clear enough in most ordinary cases whether claims are linked or are working separately.³ Any diagramming technique aspiring to adequately represent reasoning and argumentation must display this difference transparently. Some of the more sophisticated argument mapping software packages do have this capacity; these include Araucaria (Rowe, Macagno, Reed, & Walton, 2006), iLogos (Easterday, Kanarek, & Harrell, 2006), and Reason!Able (van Heuveln, 2004).

In Rationale, Analysis maps serve this purpose. Premises are shown to be linked to form a single complex Reason or Objection by having them enveloped in an appropriate color. The color can be understood as an extension of the line joining the linked premises, as a set, to the contention.

A diagramming convention unique to Rationale is the way the color fades to nothing at the bottom of the tallest premise. This convention solves a design challenge which arises because there might be arguments bearing upon any of the premises, and so we must allow for the possibility of colored lines connecting to those premises. If the color completely surrounded the premises, then the incoming lines would have to cross that color – an ugly visual effect. To avoid this, the underside of the reason is exposed to allow additional reasons or objections to be attached directly. In this way arbitrarily complex structures of argument can be created.

³ At least, it is clear enough for someone with a strong practical grasp of this distinction. And even experts can find recalcitrant cases, which are of course intimately related to the difficulty in precisely articulating the distinction.

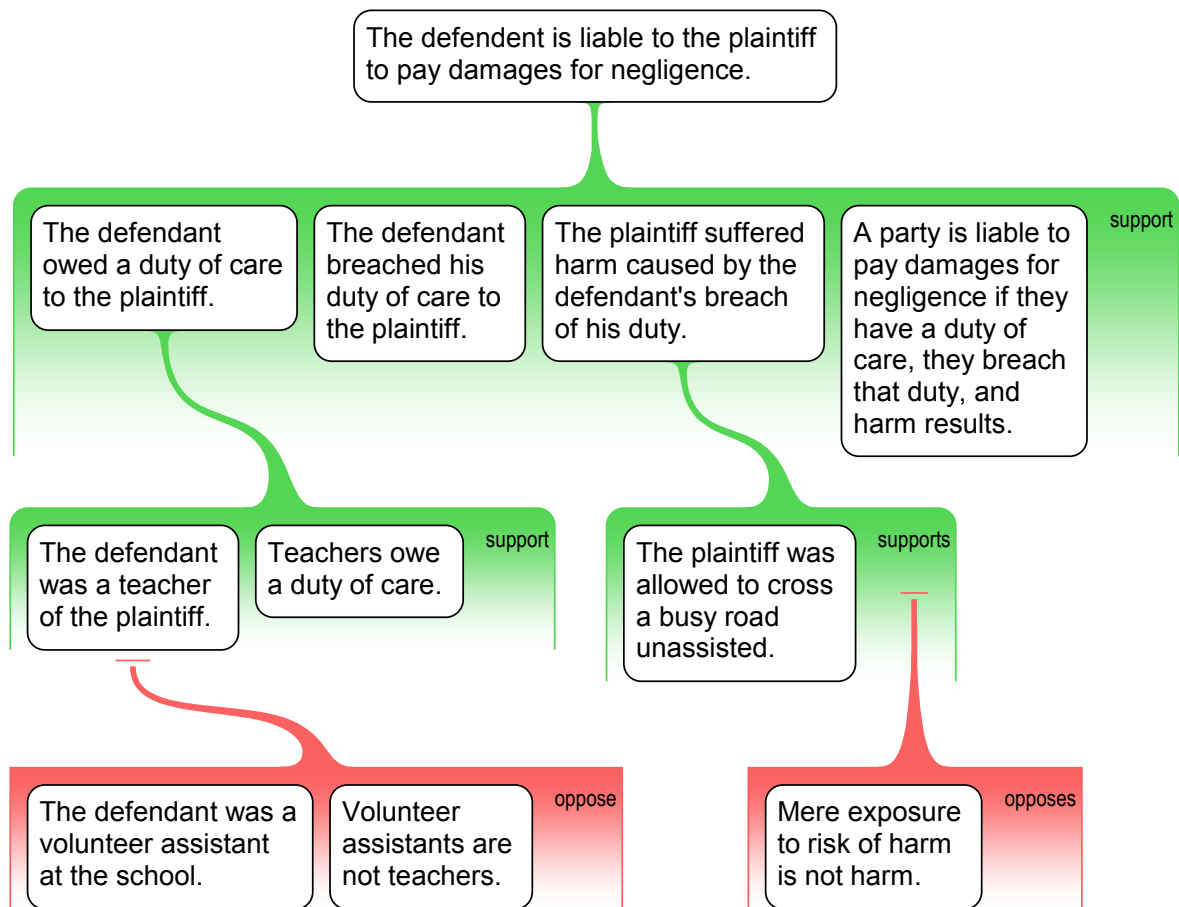


Figure 1. An Analysis map, illustrating how linked premises are shown to be constituting a single Reason or Objection. Note also the way the color fades out, allowing lower-level considerations to connect to premises without crossing visual boundaries.

The Analysis map in Figure 1 also illustrates how Rationale handles a common reasoning pattern, which we refer to as an “inference objection.” This is where somebody raises an objection to an argument, an objection countering not the stated premise but rather the relationship between that premise and the claim it is supporting (or opposing). In this example, the right-hand objection concedes that the plaintiff was allowed to cross a busy road unassisted (for the purposes of the example, a rather disturbing scenario). The force of the objection is to deny that the conceded claim proves that the plaintiff suffered harm as a result of the plaintiff’s negligence. Thus the objection connects not to the claim, but to the green area representing the

support relationship.

Multi-Map Workspace

The Rationale workspace is designed so that any number of infons and maps of various kinds can sit alongside each other on the same workspace. Infons and whole maps can be easily dragged and dropped into and out of maps and between maps. An individual infon can be free-floating, or dropped into a Grouping map as an Item, a Reasoning map as a reason, or an Analysis map as a premise. Whole maps can be transformed from one kind to another, though can result in loss of structural information.

Modifying Maps

The discussion thus far might have suggested that the main thing Rationale does is help draw diagrams of reasoning. This is indeed a central function, but an equally important role is to help us work out what the reasoning actually is. Prior to the process of mapping, we usually do not have in our minds a fully refined conception of the reasoning just waiting to be diagrammed. Rather, we generally only arrive at such a conception through an iterative process of drafting and revision. When we see reasoning laid out clearly before us, we are better able to identify gaps, obscurities, errors, etc., prompting reformulation. In cases of very complex reasoning, this can go on indefinitely.

Thus, Rationale provides considerable support for modifying diagrams quickly and easily, primarily through drag-and-drop operations, but also via right-click menus, keyboard commands, and buttons on the ribbon⁴ and other places. In the design of the software, at least as much functionality is devoted to modifying existing maps as to producing new maps - though of course there is no sharp line between these two activities. At the heart of this functionality is automated layout of diagrams. Any time the user makes a change, e.g., relocating a single box, the software automatically redraws the entire map. The algorithms supporting this automated layout are complex and are a key underpinning of the distinctive

⁴ Rationale was an early adopter of the “ribbon” interface design developed by Microsoft for Office 2007. In fact, because Microsoft made this design public well before the official release of Office 2007, it was incorporated and released in Rationale well before it officially appeared in Office 2007.

functionality and usability of Rationale.

Viewing Maps

As reasoning becomes more complex and the workspace more crowded, it becomes increasingly difficult to find and follow the reasoning of interest at any given moment. This problem is created in part by the fact that computer monitors, while steadily increasing in size, still offer limited screen "real estate" compared with what would ideally be available to display large amounts of reasoning and argumentation. It is common, for example, to be unable to see both the trees and the forest, i.e., read a particular proposition while at the same time being able to scan the full context of reasoning in which it is embedded.

To help overcome these challenges, Rationale provides various ways to modify one's view of the workspace. There is the obvious panning and zooming; an overview window in the upper right corner, so that the entire contents of the workspace is always at least schematically displayed; the ability to automatically pan and zoom to a selected map, or from map to map around the workspace; and the ability to modify the layout of any particular map to display more effectively some feature of the reasoning. The last of these is distinctive to Rationale, and an example appears in Section 4.1 below.

Evaluating Maps

The ultimate goal of most reasoning and argumentation is to establish or assess the truth (of falsity) of propositions in the light of the relevant arguments. To do this we must *evaluate* the arguments, i.e., determine how strong they are, individually and collectively. Rationale supports the user's evaluative activities by providing for evaluative judgments to be displayed as a visual overlay on argument maps using icons and variations in shading, color, and line thickness. In complex arguments, displaying evaluations can help one rapidly discern patterns of strength and weakness, in a way not possible in traditional formats.

Note that Rationale does not itself make any evaluations. The software doesn't understand the claims, cannot determine whether they are true, and cannot estimate the strength of reasons or objections. Of course it would be possible to augment Rationale with algorithms drawn from work in Artificial Intelligence. This would however be a significant departure from the original vision for Rationale (and especially its predecessor, Reason!Able), which was to be a tool

supporting the user's thinking processes, rather than being a tool which, to some degree, did the thinking for them.

3 Making People Smarter

As mentioned above, the primary purpose of Rationale is to help improve human reasoning - roughly, to "make people smarter." It does this by helping people reason more effectively, rather than doing reasoning for them. In other words, Rationale is in the game of *intelligence augmentation* (IA) as opposed to *artificial intelligence* (AI).

Broadly, there are two ways in which the software can be used to make people smarter. One is educational; the software helps build our core reasoning skills. The other is "professional;" it is improving our capacity to *apply* our reasoning skills, such as they are, more effectively in our work.

3.1 Educational Use

Rationale originated in attempts to improve the critical thinking abilities of undergraduate students. Critical thinking is widely acknowledged to be a primary goal of education, particularly higher education. General reasoning and argument skills are central to critical thinking, and so to improve critical thinking we need effective ways to enhance those skills. There is considerable evidence that this is difficult, and many have surmised that attempts to improve critical thinking have little if any effect. For example, a leader in the field wrote only a few years ago "I wish I could say that I had a method or technique [for teaching Introductory Critical Thinking] 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).

In response to this dismal situation, a research effort known as The Reason Project at the University of Melbourne took on the challenge of developing a method that was demonstrably effective while being affordable for widespread use. The guiding insight behind The Reason Project's approach was that reasoning, as a skill, should improve the same way as all other skills, i.e., through lots of dedicated practice. Such practice usually requires lots of coaching from human experts. In the case of reasoning such coaching is expensive and so rarely provided. The Reason Project attempted to sidestep this problem by providing an environment in which students could do lots of reasoning practice of a more beneficial nature than is usually

found in traditional teaching approaches. That environment was Reason!Able (van Gelder, 2000), one of the earlier argument mapping software packages to be widely available and the direct predecessor of Rationale. Reason!Able was combined with a curriculum for developing critical thinking skills in the context of a one-semester undergraduate subject. This curriculum involved lots of argument mapping exercises, primarily argument reconstruction (i.e., producing an argument map of the reasoning in some argumentative text), but also argument construction and various other exercises. The overall approach is now known as LAMP, an acronym for Lots of Argument Mapping Practice (Rider & Thomason, forthcoming).

The Reason Project was concerned, almost to the point of obsession, with determining whether the LAMP approach was actually working (i.e., is it "demonstrably effective"?). The extent to which skills were actually improving was determined by testing the students before and after undergoing LAMP training. A number of different pre/post studies, using independent tests, were performed over a five year period. These reliably indicated that students made substantial gains in critical thinking ability, with an effect size⁵ of around 0.8 (Alvarez, 2007). This may sound like a small number, but 0.8 is deemed large in a widely-used rule of thumb (Cohen, 1992). For a very rough comparison, consider that a gain in IQ of that magnitude over the same period would be equivalent to gaining one IQ point per week over 12 weeks.

In order to know whether (or the extent to which) these substantial gains were caused by the LAMP training, we need to know how much those students would have gained had they not had the LAMP training. Ideally this issue would be resolved through large trials with random assignment of "subjects" to either a "treatment" group doing LAMP training or a "control" group doing some alternative, comparable activity. However various practical barriers stand in the way of such trials in this sort of educational research. Instead, The Reason Project estimated the amount of gain attributable to the LAMP method by comparing the gain in LAMP-trained students with gains found in students in a variety of other contexts. To this end it conducted a

⁵ The statistical notion of effect size allows us to compare the results of various different studies even when they differ in important ways. The precise definition of effect size varies depending on the type of study. For the purposes of our studies, the relevant notion of effect size is the mean gain from pre-test to post-test divided by the standard deviation of performance on the test.

meta-analysis of all available studies of gains in critical thinking at university (Alvarez, 2007). The meta-analysis results enable us to have fairly high confidence in the following claims:

- Students at college or university typically improve by about 0.1SD over one semester. Here, “SD” refers to the standard deviation of performance on whatever test is being used to measure their critical thinking skills.
- Students enrolled in critical thinking subjects typically improve by about 0.3-0.4SD over one semester.
- Students undergoing LAMP training typically improve by about 0.8SD over one semester.

From these results we can provisionally infer that LAMP training is on average about twice as effective as other critical thinking subjects, and many times more effective than simply being a university student. Some caution is needed when making these inferences because it is impossible to rule out the presence of some confounding factor wholly or partly responsible for the strong gains shown by the LAMP students in our studies. However in the absence of any specific, plausible conjecture as to what the confounding factor might be, it is reasonable to conclude that LAMP is primarily responsible for the difference. And since LAMP makes extensive use of argument mapping software, software designed with the express purpose of helping students improve their critical thinking skills, it is reasonable to further conclude that using the software did in fact help those students improve their skills.

As mentioned, the LAMP studies involved students using an earlier software package, Reason!Able. Rationale is essentially the same as Reason!Able, except that it is far better quality. It is plausible that if the same studies were done with students using Rationale, they would uncover gains at least as large.

In short, there is serious empirical evidence that Rationale can be an effective aid in improving peoples' basic reasoning capacities. After practicing reasoning with Rationale, people are better at reasoning, even when not using Rationale.

3.2 Professional Use

Although Reason!Able was developed for educational use, it soon became obvious that some

such tool might be useful in helping even sophisticated thinkers handle complex reasoning. This shouldn't be surprising. It is uncontroversial that visualization can help our minds cope with complexity. Reasoning and argumentation can be highly complex. Thus, reasoning visualization tools such as Reason!Able should help people cope with reasoning and argument. Although some people did try using Reason!Able in this way, it had many usability problems. Hence there was a need for a package developed from the outset to support people engaged in "real world" reasoning tasks, particularly in workplace contexts. Rationale was conceived as being comparable to, and sitting alongside, other office productivity applications. Whereas a word processor handles text, and a spreadsheet is a number processor, Rationale is an argument processor.

Does Rationale in fact help professionals reason better "on the job"? At this stage all we have is anecdotal evidence. For example, lawyers at one of Australia's largest firms were representing one party in one of the larger and more controversial Australian competition law cases for many years. The case concerned whether one multinational mining company had the right to deny another smaller mining company access to the railway line it had built for transporting ore from its inland mines to a port on the remote coast of Western Australia. The case turned on the interpretation of a section of the Trade Practices Act, which says, roughly, that one party is entitled to use a service provided by a facility constructed by another *unless* that service is a production process. Unfortunately key notions, such as *production process*, were obscure, and a forbiddingly complex tangle of argumentation had enveloped the dispute. At the same time, given the enormous stakes involved, the lawyers assigned to the case had to be sure they had a crystal-clear understanding of the arguments and could mount the strongest possible case. They used Rationale to help lay out the arguments in a transparent, ordered fashion, and to develop new arguments (Drummond, 2006). When the final decision was handed down, it was in their client's favor.

There are other examples - policy analysts using Rationale to develop and defend policy, intelligence analysts using it to test and support their findings, managers using it for strategic analysis. In all such cases, it is intrinsically difficult to determine the degree to which Rationale use leads to better reasoning and successful outcomes, but the larger the body of anecdotal evidence, the more plausible it becomes that there is a real effect. Ideally, there would be rigorous empirical studies aimed at teasing out the extent to which reasoning "on the job" is in

fact improved, and some such studies are now being launched.

4 Why does it work?

Assuming that Rationale really does, or at least can, make people smarter, we might ask why this is so. What is it about Rationale, or argument mapping software more generally, which helps us reason more effectively? Because argument mapping is a new phenomenon, there has so far been little research in this area. We are only gradually developing an understanding of the relevant issues. At least three main themes are emerging: usability, complementarity, and semi-formality. These are not three independent explanations; they are better thought of as overlapping "takes" on how or why Rationale achieves its intended effect.

4.1 Usability

In a nutshell, the first claim is this: a tool like Rationale improves reasoning because it is highly *usable* for reasoning activities, or at least more usable than relevant alternatives.

This is not simply the assertion that the software is "user friendly," which usually means that the software is attractive to, and easily used by, naive users. Rather, the technical notion of usability concerns the degree to which a tool or system enables standard users to conduct their activities or achieve their goals effectively and efficiently, and perhaps also with some measure of pleasure or satisfaction.⁶ A usable tool might not be very useful for naive users. A good analogy here is windsurfing. There are basically two kinds of windsurfing boards. Beginner boards are large, stable, and support a person even when not moving. They are very "friendly" to windsurfing naïfs. Regular or advanced boards are smaller, less stable and more nimble; they are very difficult for beginners but support a far better windsurfing experience for those who are competent. The assertion that Rationale is highly usable means primarily that, like the advanced windsurfing board, it enables people who are competent in reasoning and the use of the tool to engage in reasoning activities more effectively, efficiently and satisfyingly.

The claim that Rationale improves reasoning because it is highly usable plays out differently in

⁶ The International Organization for Standardization, in its document ISO 9241-11 (1998), defines usability as "The extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use."

each of the two main contexts of use. In the educational context, Rationale's usability for reasoning helps improve reasoning *skills* by enabling a student to do more practice, and practice of a better kind, than they can do using traditional techniques. Just as you can become a much more skilful windsurfer through lots of practice on an advanced board, so you can become a better reasoner through lots of use of a tool like Rationale, even if it takes some training to get "up to speed." There is an important disanalogy however. In windsurfing, the skill you acquire can only be deployed on a suitable board; whereas in reasoning, the skills you acquire are more generic and transferable, and can be deployed even without the software tool which enabled the development of those skills.

In the professional context, the claim that Rationale improves reasoning performance because it is highly usable is a tautology, since to be usable is, by definition, to enable better performance. However invoking the notion of usability still helps because it points in a certain direction. Exploring the issue from a usability perspective can help us better understand why and how a software package like Rationale improves performance.

When we claim that a tool like Rationale is highly usable, we are not measuring Rationale's usability on some independent, objective scale. Rather, we are saying that it is significantly more usable than relevant alternatives. Thus to make sense of the claim, we need some understanding of what the relevant alternatives are. What tool or tools do we standardly use to help us engage in informal reasoning or argumentation? If engaged in a complex debate over, say, carbon trading, or the war in Iraq, what do we use to help us organize and evaluate the arguments?

The answer is that, overwhelmingly, we use prose. By this I mean that we articulate our arguments using sentences organized sequentially on a page or pages, and using various means such as indicator words ("therefore" etc.), paragraphs, indentation and dot points to help illuminate how the parts of the arguments hang together. We use prose as a tool to help us *develop* arguments, as when we figure out what our argument is through producing and editing drafts; and to *present* arguments to others and even ourselves.

Argumentative prose can be considered a very generic or abstract kind of tool. Our use of prose is supported in turn by various more concrete technologies. For example we can use pen and paper, or we can use a word processor running on standard personal computing hardware.

The support technologies have changed substantially over time but the way we use those technologies to engage in reasoning and argumentation has remained essentially constant from the time of the ancient Greek philosophers through to the present day, which is why the works of Plato and Aristotle can still be part of the canon studied even by undergraduate philosophy students.

A tool can be very widely *used* even if it is not particularly *usable*. There are innumerable examples; thus, not so long ago, fountain pens were widely used for writing on paper, but they are less usable than ball-point pens, even if the result is sometimes more aesthetically pleasing. In the case of prose, people are generally so accustomed to using this tool, so ignorant of any alternative, and so blind even to the idea that prose is the tool they are using, that they fail to realize what its usability problems are or even that it has usability problems. Any deficiencies in a person's reasoning, or presentation of that reasoning, are attributed to deficiencies in their education or their intellectual capacities rather than being traced, at least in part, to inadequacies in the tools.

An "argument processor" such as Rationale, based not around prose but around specially-designed diagrams, is an alternative to prose. Importantly, it is an alternative that was developed with the deliberate goal of being more usable than prose, so it would not be surprising if it turned out in fact to be more usable. The interesting part is how it manages to be so.

First, any contemporary software tool is able to take advantage of the wisdom accumulated in decades of research into how interactive tools should be designed so as to best support our activities, particularly our cognitive activities. The lessons learned from this research are increasingly encapsulated in authoritative sources (e.g., (Raskin, 2000)) and embedded into the tools and conventions used to develop contemporary applications. They concern diverse issues in the design of a tool meant to maximize performance and an experience of "flow": when and how to use (or avoid) "modes" or dialog boxes, how to use size and color; how to align behavior with user's mental models; etc. Rationale is in the fortunate position of being able to exploit this accumulated general wisdom and apply it, almost "off the shelf," to the case of a tool to support reasoning activities.

Second, an argument mapping package can take advantage of a wider range of basic

representational resources. Consider some typical argumentative prose - for example, an opinion piece in a major newspaper. What means does the author use to help convey to you, the reader, how the various key claims hang together as an argument? Most obviously, the author might use explicit verbal indicators - phrases such as "This is because..." and "Hence...". Other tactics include word and sentence ordering; paragraph breaks; and subtle cues based in the meaning of terms and the context in which the piece is written. Upon reflection this is a remarkably meager set; it almost willfully ignores a range of other resources available not just to the contemporary computer user but even to a child with just paper and a bag of colored pencils. These resources include:

- symbolic argument structure markers, such as the philosophers' standard "P1, P2, ... C"
- color
- shapes
- lines or arrows
- position in space

If your goal is to produce displays of reasoning which maximize comprehension, manipulation, evaluation and communication, why wouldn't you take advantage of such cheap but powerful visual aids? If we can break the shackles of convention and habit, we are free to exploit any resources which in practice can aid the process of reasoning. Argument mapping software helps itself to these resources, thereby gaining an "unfair" advantage over traditional argumentative prose.

Third, a tool like Rationale is adapted or tuned to the unique demands of reasoning and argumentation activities. Prose, and supporting technologies such as word-processors, are generic; they can be used for reasoning but are not designed specifically for such use. A purpose-built tool provides distinctive ways of representing reasoning structures and "affords" appropriate kinds of operations on those representations. Its design is constrained by the nature of reasoning activities, while not being distorted or diluted by the need to support activities other than reasoning.

As an example, consider again the feature mentioned in Section 2 (Viewing Maps), the ability to modify the layout of a map to reveal some aspect of the reasoning more effectively. One such

aspect is the connection between the main point at issue and some proposition buried deep in the reasoning. For example, in the case of the arguments surrounding the assassination of JFK, it has been pointed out that the Zapruder film seems to show JFK's head moving backwards at the time of the final, fatal shot. But how exactly does this seemingly minute detail relate to the main issue so furiously debated, that there was a conspiracy to kill JFK? We would like to see the chain of inference from that claim (or if we prefer, any other claim in the complex debate) and main contention. In Rationale, this is done via the Show Ancestors command, which lays out the map with the chain of inference from the selected claim to the main contention displayed in a straight line, with other parts of the argument moved aside to accommodate that line.

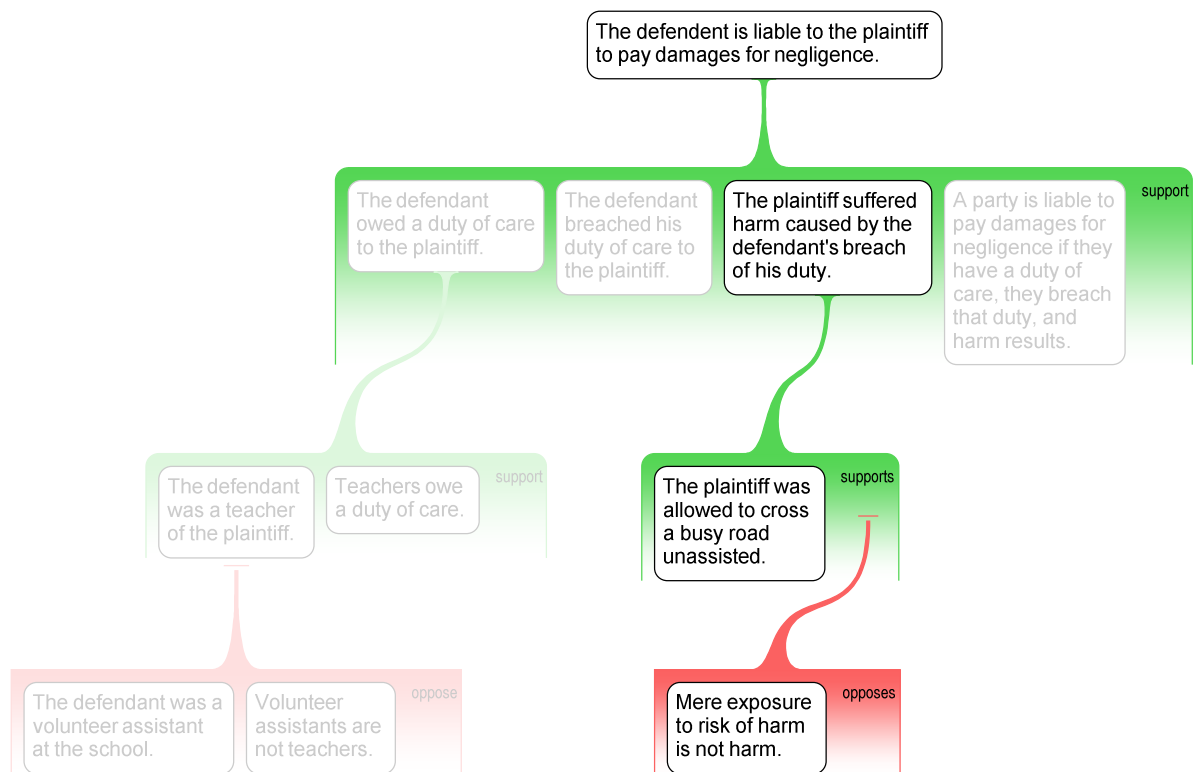


Figure 2. Effect of the Show Path command, highlighting the chain of inference connecting a claim with the main contention.

Another aspect of this layout is worth mentioning. Peter Tillers has described the importance,

in comprehending complex reasoning, of being able to focus attention on a particular part without losing awareness of the larger structure in which it is embedded (Tillers, 2005). He recommended that argument visualisation software be designed so that surrounding arguments can be visible but de-emphasized, supporting a peripheral awareness of those arguments. As illustrated in Figure 2, Rationale handles this automatically for certain layouts by making surrounding arguments semi-transparent, and allows the user to adjust the level of transparency.

Rationale makes possible this kind of rapid reconfiguring because it helps people understand complex reasoning, and the whole point of Rationale is to help people with reasoning. Conversely, this kind of redisplay is impossible when using technologies, such as traditional argumentative prose, which are not designed specifically for the task.

Current packages such as Rationale are just stages in an ongoing process of redesign, the ultimate goal of which is a tool so usable that it becomes like an invisible extension of our cognitive apparatus. Just as we are not aware of our own brains when we are thinking, but are aware of what our brains are helping us think about, so an ideal argument mapping package would be like the blind man's cane, something "through which" our minds engage in complex deliberation, conscious only of the reasoning itself, not of any issue or difficulty in dealing with the tool. No package has reached that goal yet, and perhaps no package ever will; but the best of the current generation are at least, by design, getting significantly closer.

4.2 Complementation

A second theme in understanding how a tool such as Rationale can make us smarter is *complementation*: the tool complements our minds' natural strengths and weaknesses.

Our reasoning abilities are a function of our basic cognitive capacities, which depend in turn on our brains. Those brains are the result of a long, accidental, accretive and incomplete process of evolution. They were "designed" not for the sublime arts of logic and disputation but to help us to survive and propagate in the physical and social environments of our ancestors going back long before the emergence of homo sapiens. Our reasoning abilities are a late acquisition, due not to any newly-crafted reasoning module but rather to our learning how to exploit in a new way cognitive capacities already wired in for other purposes. Those capacities include some

remarkably powerful and useful functions, but do not provide everything that might be needed for flawless, general-purpose reasoning.

In light of this, a sensible approach to improving human reasoning is to provide tools which, on one hand, provide ways to bypass or make up for the deficiencies or limitations of our innate capacities, while on the other, taking advantage of their distinctive strengths. Rationale does both of these.

The most important respect in which a tool like Rationale compensates for inherent cognitive limitations is by augmenting our short-term memory (STM). The need for such augmentation can be summarized in three points. First, reasoning and argumentation can become very complex. Second, as we standardly do things, that complexity places huge obligations on our memories. However, third, our innate STM is quite limited - much more so than people usually realize.

The standard cliché about human STM is that it can hold "7+/-2" items. Thus 8 digit phone numbers are quite a bit harder to remember than 7 digit ones. The 7+/-2 figure originated in famous research in the 1950s (Miller, 1956). However the figure should be treated with caution. The original research concerned our capacity to recall random sequences of meaningless items. More recent research suggests that for such sequences, Miller's figure may be an overestimate (Cowan, 2000). At the same time, STM can be increased when items are "chunked" or meaningfully grouped. And with suitable training, people are able to perform remarkable feats of short-term memorization, such as remembering dozens of random digits (Wilding & Valentine, 2006).

Nevertheless, it is clear that, in normal circumstances, human STM has quite a low ceiling. This can be easily illustrated by asking two people to play tic-tac-toe without using pen and paper or any equivalent - i.e., to hold the state of the game in their head rather than on a board. Most people find that completing the game is a fun challenge, in part because it is hard enough just keeping track of the moves already made, let alone plotting clever new moves.

Meanwhile it is obvious that arguments can be very complex - far more so than a mere game of tic-tac-toe. For example the case put forward by Jim Garrison in the trial scene of Oliver Stone's movie JFK consists of dozens of pieces of evidence woven into an intricate web. Even this case (let alone the galaxies of arguments and responses in the larger "Who killed JFK?" debate) is

more complex than can be held, organized and evaluated purely in the head by anyone other than, perhaps, an idiot savant.

This point holds true even in the mundane territory of our everyday disputes or academic altercations. The issues and arguments are often if not always larger than our unaided minds alone can easily embrace; and retaining even some of that material is effortful and prone to loss, confusion and confabulation.

Yet we do presume to think through such complex cases, and so by practical necessity we make use of external *aides de memoir*; for example, recording and organizing our thoughts in notes, essays or books - i.e., in some form or other of prose. Similarly, a central function of an argument mapping program such as Rationale is to function as an external "memory", or memory-extension, for reasoning. Maintaining stable structured representations of arguments or debates of effectively unlimited complexity is trivial for an appropriately-programmed computer. Thus, when using such a package to help us think our way through a set of arguments, we are taking advantage of a great strength of computers to compensate for an inherent weakness in our own capacities.

For external representations of reasoning to be useful in our thinking, we must be able to interact with them fluidly. On one hand, we must be able to create and modify the representations easily; this was touched upon above. On the other, we need high-bandwidth access to the information contained in those representations. Ideally, we'd be able to "read" the representations faster than we could think about what they are representing. Here, argument maps are designed to exploit some remarkable strengths of our standard cognitive equipment. It was pointed out above that an argument mapping software package can be more usable as a tool for reasoning because it exploits basic representational resources which are neglected by typical argumentative prose - resources such as color, shape, line, and position in space.

The use of these resources is a great advantage because our "hard wired" mechanisms for visual cognition are designed to process information coded in these basic dimensions with extraordinary efficiency. When you look outside the window and see a tree, your perceptual system accepts and processes a vast amount of basic visual information and reliably delivers a correct high-level interpretation in a fraction of a second and with no perceptible effort on your part. This is an example of what psychologists often call "pre-attentive processing" -

information being taken up and utilized so fast that you didn't even have time to shift your attention to it (Ware, 2004).

Color, shape, line, and position in space are all aspects of a scene which are, or can be, pre-attentively processed. Any tool for reasoning which wants to optimize the transfer of information from external representations to central cognitive processes ought to exploit pre-attentive processing as much as it is effectively able. Argument mapping software takes this approach, with the result that representations of complex reasoning can be accessed, and hence utilized in thinking, more quickly and easily than in standard prose formats.

The most obvious example of this is the use of color to code for polarity, i.e., whether one proposition is (taken to be) supporting another proposition ("positive") or opposing it ("negative"). In standard argumentative prose, polarity is something which must be "computed" through slow, effortful and error-prone high-level interpretative processes. Consider for example this passage:

Thus, *apres le postmoderne*, realism continues to be of the utmost significance. First, because a great portion of high and popular culture texts continue to exemplify classic realism. But perhaps even more importantly because realist ontologies condition the reception even of modernist and postmodernist cultural forms. To be sure, these popular ontologies presuppose a different, and far more contingent, even cynical, vision of what is real than did their nineteenth century counterparts. And to be sure, much of turn-of-the-twenty-first-century realist culture will be of a qualitatively more 'brute', disharmonic nature than its predecessors. The new realist culture may indeed be cast in a mould that is neither semiotic, nor mimetic, but starkly indexical in coloration. Yet, this said, realist cultural forms will persist. And those who will increasingly consume modernist and postmodernist culture will do so for realist reasons. Critics of the twenty-first century, then, will be well advised once again to take realism seriously. (Abercrombie, Lash, & Longhurst, 1992)

This is a quite ordinary piece of argumentative academic prose. It is easy enough to follow the gist, at least if you are familiar with words such as "ontologies" and "disharmonic." However identifying the exact logical structure is quite intellectually taxing. You must interpret the

sentences, and think carefully about how they relate to each other. In our experience, even highly educated people struggle to identify how the pieces fit together, and they tend to arrive at wildly different interpretations.

If however the core argument is presented in a standard argument mapping format, then, assuming you are well-versed in the relevant conventions, you will pre-attentively process the colours, shapes, positions and lines and and rapidly and easily see the basic structure of the reasoning. In other words information about evidential structure is being conveyed in manner which is extremely fast, almost effortless, and unambiguous - much like seeing the tree outside the window.

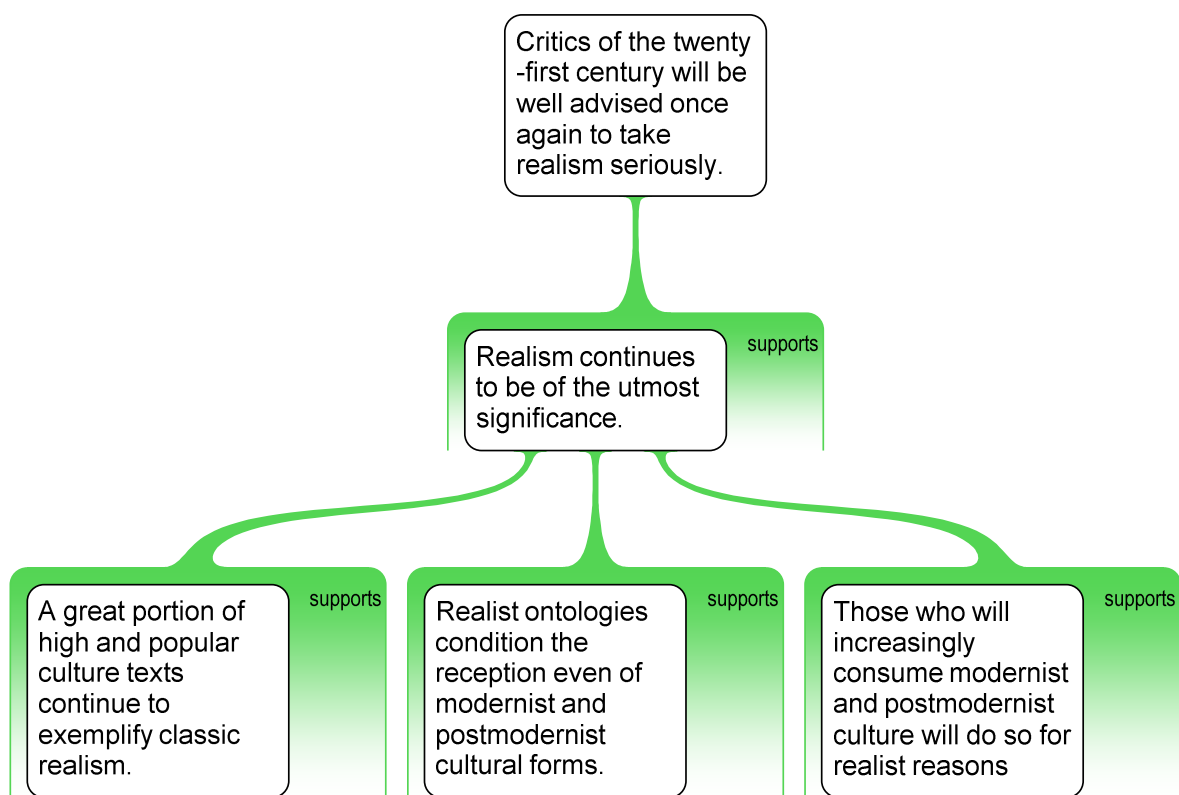


Figure 3. A simple Rationale Reasoning map, illustrating how pre-attentive processing can help convey argument structure.

Thus, an argument mapping software package is exploiting an impressive strength of the

human mind, namely its extraordinary facility in processing certain kinds of basic visual information.

More profoundly, among the greatest strengths of the human mind are its abilities to comprehend natural language and evidential relationships. Despite the valiant efforts of computer scientists over some decades, computers still lack anything seriously resembling general human intelligence, notwithstanding world-champion chess programs and the like. Thus a software package designed to improve human reasoning must still rely on our minds to do core "heavy lifting".

4.3 Semi-formality

A third theme in explaining how using an argument mapping tool such as Rationale can improve reasoning centers on the notion of *semi-formality*. In a nutshell, the idea is this: human thinking is typically informal. In certain areas, such as general reasoning and argumentation, it can be improved by moving into a semi-formal, rather than formal, mode. Argument mapping is reasoning conducted in a semi-formal mode.

The contrast between formality and informality, in the sense relevant here, is familiar enough. Modes, domains, systems, languages or whatever count as formal to the extent they have characteristics such as:

- a finite set of basic symbols or primitives
- are digital (Haugeland, 1985)
- are governed by strict rules or algorithms
- have strictly-defined concepts or "semantics"
- achieve complexity through combination
- support effective procedures

Mathematics, formal logic, programming languages, artificial intelligence systems, and many games such as chess are all formal in this sense. Natural language, conversation, politics, and humor on the other hand are informal, even if they exhibit a few of the features of formality to some extent.

Now human reasoning and argumentation are standardly informal. While there are primitives (words, or propositions), meanings, and principles or norms, these are not defined with the kind of precision, rigor or reliability one finds in formal modes such as mathematics or chess. In this respect, human reasoning is often reflecting the nature of the domains or issues about which the reasoning is being conducted. For example, politics is an inherently informal domain, and this informality is reflected in the nature of our default intellectual tools for thinking about and debating over political topics.

Noting on one hand the problems associated with ordinary human reasoning and argumentation, and on the other the impressive achievements of their formal counterparts, the temptation has often been to recommend shifting human reasoning into a formal gear. Thus introductory logic textbooks are usually dominated by elementary formal logic (Aristotelian syllogisms, propositional logic, predicate calculus), making the assumption that people would reason more effectively if they replaced their instinctually informal thinking habits with logical formulae and proofs. This tendency culminates in the aspiration of mainstream artificial intelligence to recreate human-grade intelligence in the formal medium of digital computation.

Unfortunately this generally doesn't work. AI research has discovered that it is extremely difficult, and perhaps impossible, to engineer formal systems capable of reliably making the mundane inferences underlying everyday conversation or humor, let alone engaging complex argumentation of the kind found, for example, in legal practice (Dreyfus, 1992). Conversely, when people struggle with everyday reasoning and argumentation, they are generally not helped by attempting to translate their premises into some logical calculus and draw inferences by reference to its official rules.⁷ Indeed for innumerable commonplace reasoning challenges, formal techniques are so hopelessly impractical that recommending their use seems like a kind

⁷ Introducing a symposium entitled "The Role of Formal Logic in the Evaluation of Argumentation in Natural Languages," eminent logician and philosopher Y. Bar Hillel said: "It is ... of vital importance to get better insights into the nature of argumentation in natural languages, and I challenge anybody here to show me a serious piece of argumentation in natural languages that has been successfully evaluated as to its validity with the help of formal logic. I regard this fact as one of the greatest scandals of human existence." (Bar-Hillel & others, 1969) p.256. (Presumably that last sentence was humorous.) In the symposium, none of the other eminent logicians and philosophers met the challenge.

of sophisticated sadism, another grim manifestation of Francis Bacon's Idols of the Theater.⁸

There are of course notable exceptions to the claim that human reasoning is not improved through the adoption of formal techniques. Clearly, for example, the proper use of statistical methods can help us draw better conclusions about subtle correlations and causal relationships. Formal modes of thought certainly have broad and important areas of application. The point being made here is that *in general* our practices of informal reasoning and argumentation cannot be enhanced by transposing into a formal key.

It may be, however, that some degree of formality can be helpful even where full-blown formal modes such as mathematics, formal logic and computation get no useful traction. The distinction between formality and informality is not a binary opposition; one might say, the distinction is not itself a formal matter. Rather, there is a spectrum of cases depending on which of the characteristics of formality are adopted and how those characteristics are manifested.

In the case of reasoning and argumentation, there have been many contexts in which these practices have been made *more* formal than they would normally be. Consider the medieval theory of disputation, which presented a sophisticated framework of rules governing moves in argumentation (Spade, 1988); or for a more mundane example, consider contemporary high-school debating or "forensics." These sorts of elaborations of our ordinary practices are aimed, at least in part, at improving those practices; i.e., at enabling us to reason and argue more effectively. Implicitly they are proposing that the optimal mode for human reasoning is not the informal or the formal, but rather something intermediate, a compromise or perhaps a "best of both worlds" scenario.

The conjecture, then, is that for general reasoning and argumentation there is a "sweet spot" somewhere on the spectrum between ordinary informal practices at one end, with their sloppiness and disorder, and purely formal techniques at the other, with their rigidity and limited range of application.

⁸ "Lastly, there are idols which have immigrated into men's minds from the various dogmas of philosophies, and also from wrong laws of demonstration. These I call Idols of the Theater; because in my judgment all the received systems are but so many stage-plays, representing worlds of their own creation after an unreal and scenic fashion." From (Bacon, 1905 (1620))

Argument mapping pushes reasoning in a formal direction by asking for a relatively high level of explicitness and rigor in articulation. It generally requires that:

- Claims be rendered discretely in full grammatical sentences
- Material not directly involved in the reasoning be stripped away
- The central contention or point in dispute be identified as such
- All direct evidential links between claims be specified
- Unstated claims or "assumptions" be identified and explicitly stated

In Rationale, these activities are scaffolded by providing a graphical format or "syntax" for the articulation of reasoning, with a limited set of "primitives" (claims, reasons, objections etc.) and strict rules about how those primitives can be combined.

Beyond the basic constraints, there are a number of principles of good mapping which cannot be specified as strict, universally-applicable rules. For example there is the principle of abstraction: generally, in a well-developed argument map, claims "higher" in the map (i.e., closer to the main contention) should be more abstract, and claims lower down should be more particular or concrete; and claims at a given level should be approximately the same in their degree of abstraction (Minto, 1995). While this principle is easy enough to state and understand, and with some practice is easy enough to apply in most cases, it involves inherently vague notions and is subject to a wide variety of exceptions, such that any attempt to articulate in a fully precise way what the principle is and how it applies sinks in a quagmire of uncertainties and exceptions.

The principle of abstraction and its ilk depend fundamentally for their successful application on intuitive human judgments based on experience and practice. To the best of our knowledge it is impossible to cash them out as formal rules capable of mechanistic implementation. Thus, argument mapping has irreducibly informal dimensions even as it makes reasoning activities *more* formal than they would normally be.

So argument mapping is semi-formal; it introduces aspects of formality while acknowledging its limits and retaining essentially informal dimensions. The idea is that disciplining reasoning practices to observe this degree of formality is the most feasible way to, if not eliminate, at least substantially mitigate the typical failures or difficulties standing in the way of good reasoning,

argumentation and deliberation, such as:

1. not making reasoning fully explicit, including in particular the failure to articulate key assumptions
2. not coping with the complexity of "real world" debates
3. not applying relevant principles of good reasoning
4. not achieving common understanding among participants

Argument mapping helps mitigate these problems, which is why we can be optimistic that argument mapping practices, if widely and properly adopted, can lead to better reasoning. It should be clear from what has been said already how argument mapping confronts the first two. Argument mapping helps address the third because, once reasoning has been made fully explicit and more easily comprehensible, it becomes easier to apply relevant principles of good reasoning. The fourth is addressed if argument maps are used not just by an individual, but in a collaborative fashion, where everyone is looking at, discussing and modifying the same map.

What is the optimal level of formality to introduce into our reasoning practices? Where precisely is the "semi-formal sweet spot"? We don't yet have definitive answers to these questions. Argument mapping as supported by a package such as Rationale constitutes one take on where the sweet spot is. It may have erred in one direction or the other. However evidence of the kind discussed in Section 3 above suggest that Rationale's take could well be approximately correct.

5 Conclusion - Rationale and Legal Reasoning

There is a very general trend, at this point in history, towards using visual or diagrammatic forms of expression. The law, traditionally one of the most discursive of all disciplines, is following this trend, if somewhat reluctantly. Lawyers are increasingly looking for ways to show rather than tell. Argument mapping, using a tool such as Rationale, tries to answer this need in one specific, but central, aspect of the law – the expression of complex reasoning and argumentation.

Discussions with practicing lawyers suggests three main ways in which software-supported argument mapping may potentially improve performance and, ultimately, their firms' profitability.

- **Quality.** It may improve the quality of reasoning itself, allowing lawyers to make correct inferences, devise strong arguments and counter-arguments, draw correct conclusions, and resolve disputes more effectively.
- **Efficiency.** It may improve the efficiency with which legal tasks are conducted. For example, a common task in most firms is producing letters of advice. This usually involves a junior lawyer working under the direction of a partner. At some point, the partner will review a draft letter which presents the main points of advice with supporting arguments. The process may be more efficient if, prior to the drafting of letters, the junior lawyer and the partner work over diagrams of the core reasoning.
- **Communication.** Lawyers often need to communicate complex reasoning - in court, to clients, to other lawyers, and indeed to themselves. Notoriously, their attempts to do this, in the traditional prose format, are frequently obscure if not utterly incomprehensible. Argument mapping may help convey their reasoning more directly. This would not work for all audiences; ageing judges and trial attorneys are likely to be particularly unreceptive. However many people do find diagrams more accessible than dense prose. One lawyer has told us how some clients, when presented with complex written advice supplemented with an argument map, barely glance at the prose but dive into the details of the argument on the map.

It is an old cliché that a picture is worth a thousand words; and it is often pointed out that, conversely, words can express what pictures cannot (Fodor, 1975). These theses contrast words and pictures as if they are incompatible alternatives. But argument maps blend the visual and the verbal into integrated displays exploiting the strengths of both; they are an example of *visual language* (Horn, 1999). Since we cannot easily generate such displays in our minds' eye, we need them laid out before us; and hence we need tools which can easily and efficiently do the laying out. Rationale is such a tool.

Acknowledgments

Preparation of this paper was supported in part by the Jacob Burns Fund, Cardozo School of Law, Yeshiva University. The paper was improved in the light of helpful comments from an anonymous reviewer.

References

- Abercrombie, N., Lash, S., & Longhurst, B. (1992). Popular Representation: Recasting realism. In S. Lash & J. Friedman (Eds.), *Modernity & Identity* (pp. 115-140). Oxford: Blackwell.
- Alvarez, C. M. (2007). *Does Philosophy Improve Critical Thinking Skills*. MA Thesis, Department of Philosophy, University of Melbourne.
- Bacon, F. (1905 [1620]). *Novum Organum* (R. S. Ellis, J., Trans.). London: Routledge.
- Bar-Hillel, Y., & others. (1969). Formal logic and natural languages: A symposium. *Foundations of Language*, 5, 256-284.
- Buzan, T., & Buzan, B. (1994). *The Mind Map Book*. New York: Dutton.
- Cohen, J. (1992). A primer on power. *Psychological Bulletin*, 112, 155-159.
- Cowan, N. (2000). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24, 87-185.
- Dreyfus, H. L. (1992). *What Computers **Still** Can't Do: A Critique of Artificial Reason*. Cambridge MA: The MIT Press.
- Drummond, M. (2006, 20 October). The jury's still out on a new way of thinking. *Financial Review*.

- Easterday, M., Kanarek, J., & Harrell, M. (2006). Design Requirements of Argument Mapping Software for Teaching Deliberation. In T. Davies & B. Noveck (Eds.), *Online Deliberation: Design, Research, and Practice*: CSLI Publications.
- Fodor, J. A. (1975). *The Language of Thought*. New York: Crowell.
- Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. Cambridge MA: MIT Press.
- Horn, R. (1999). *Visual Language: Global Communication for the 21st Century*. Bainbridge Island, WA: Macrovu Inc.
- 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.
- Miller, G. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81-97.
- Minto, B. (1995). *The Pyramid Principle: Logic in Writing and Thinking*: Pitman Publishing.
- Norman, D. A. (1994). *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Reading MA: Addison Wesley.
- Raskin, J. (2000). *The Humane Interface: New Directions for Designing Interactive Systems*. Boston: Addison Wesley Professional.
- Rider, Y., & Thomason, N. (forthcoming). Cognitive and pedagogical benefits of argument mapping: L.A.M.P. guides the way to better thinking. In A. L. P. Okada & S. Buckingham Shum & T. Sherborne (Eds.), *Knowledge Cartography - software tools and mapping techniques*: Springer.
- Rowe, G., Macagno, F., Reed, C., & Walton, D. N. (2006). Araucaria as a Tool for Diagramming Arguments in Teaching and Studying Philosophy. *Teaching Philosophy*, 29, 111-124.
- Snoeck Henkemans, A. F. (2000). State-of-the-Art: The Structure of Argumentation. *Argumentation*, 14, 447-473.
- Spade, P. (1988). *Lies, Language and Logic in the Late Middle Ages*. London: Variorum Reprints.
- Tillers, P. (2005). Picturing factual inference in legal settings. In R. Wittman (Ed.), *Gerechtigkeitwissenschaft: Kolloquium aus Anlass des 70. Geburtstages von Lothar Philipps*. Berlin.
- van Gelder, T. J. (2000). Learning to reason: A Reason!-Able approach. In C. Davis & T. J. van Gelder & R. Wales (Eds.), *Cognitive Science in Australia, 2000: Proceedings of the Fifth Australasian Cognitive Science Society Conference*. Adelaide: Causal.
- van Heuveln, B. (2004). Reason!Able: An Argument Diagramming Software Package. *Teaching*

Philosophy, 27, 167-172.

Walton, D. N. (2000). Problems and useful techniques in teaching argumentation, informal logic and critical thinking. *Informal Logic*, 20, *Teaching Supplement*, 35-39.

Ware, C. (2004). *Information Visualization: Perception for Design* (2nd ed.). San Francisco: Morgan Kaufmann.

Wilding, J. M., & Valentine, E. R. (2006). Exceptional memory. In K. A. Ericsson & N. Charness & P. J. Feltovich & R. R. Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 539-552). Cambridge: Cambridge University Press.