A Computational Architecture for Heterogeneous Reasoning

Jon Barwise
Departments of Mathematics, Philosophy and Computer Science
Indiana University

John Etchemendy
Department of Philosophy
Stanford University

Abstract

Reasoning, problem solving, indeed the general process of acquiring knowledge, is not an isolated, homogenous affair involving a one agent using a single form of representation, but more typically a complicated, collaborative, heterogeneous activity. This paper describes an effort to expand our understanding of such reasoning and to develop tools to enable individuals and groups to use computers more effectively in practical problem-solving tasks.

Natural deduction and problem solving

A recent article in the *New York Times* reported the discovery of mass in the neutrino by a team of 120 scientists from 23 research institutions. The discovery involved the design and construction of a massive experiment involving a tank inside a deep zinc mine, filled with 12.5 million gallons of water, and equipped with specially designed light amplifiers covering the inside of the tank. Using this set-up as a neutrino detector to compare "flavors" of neutrinos coming directly from the atmosphere versus those coming through the earth, the scientists were able to determine that some neutrinos changed flavor in passing through the earth. The discovery also had a logical element, with mass being the only plausible explanation for the observations consistent with quantum theory that could not be ruled out in one way or another.

As this example illustrates, large-scale collaborative projects involve many people reasoning toward the solution of a common problem over an extended period of time. The design and construction of a product, for example, whether a scientific apparatus, a building, or a complex hardware or software system, often involves a multi-disciplinary team of clients and engineers working toward a common goal.

Such distributed reasoning projects frequently yield less than optimal results, or even flounder. For example, the *Times* article reported a similar effort that ran over budget and ended before the apparatus was even complete. One reason for less than optimal results of such large scale efforts is the lack of techniques for structuring and recording the complex rationales that guide decisions made over the course of the project (Fruchter & Krawinkler 1995; Fruchter 1996; Fruchter, Reiner, Leifer & Toye 1998). Even when a successful outcome is achieved, later modification of the product or reuse of design components is often hampered by the absence of an easily understood record of these rationales (Parnas & Clements 1986; Conklin & Begeman 1989; Chandrasekaran, Goel & Iwasaki 1993; Bailin 1997).

We believe that a crucial piece of intellectual work is missing, one that would greatly contribute to the success of such collaborative problem solving. That work is a robust, theoretically informed computational framework for producing an accessible and comprehensible project memory, one that is structured as a rational reconstruction of the (individual and collective) reasoning that led to the ultimate design. (Current tools for capturing project memory typically yield records structured only by temporal information recorded in filenames and time stamps.)

A theoretically informed solution to this problem requires an understanding of the nature and structure of reasoning, or more precisely, rational justification. We propose that the overall structure of rational justification described by the theory of *natural deduction* (e.g. Gentzen 1935/1969; Fitch 1952; Prawitz 1965, 1973) is a reasonable model of the large-scale structure of design justification and other forms of problem solving.

In making this claim, we mean natural deduction in its broadest sense, where proofs are seen as recursively structured rationales that display the overall case structure of reasoning and permit the nesting of proofs within proofs. The claim is to be distinguished from the patently false claim that reasoning steps in the design and problem solving process take the form of introduction and elimination rules tied to specific items of logical vocabulary.[1]

Heterogeneous Reasoning and *Hyperproof*

Even when natural deduction is construed broadly, the realities of collaborative problem-solving, as embodied in the distributed design process, for example, raise problems that preclude the ready application of this model or, for that matter, any established logical theory. Our own work addresses three important problems that arise in providing tools and methodologies for recording and, where possible, verifying, a structured rationale: the heterogeneity of representation, the heterogeneity of rationale, and the heterogeneity of goals.

- **Heterogeneity of representation.** Most reasoning problems require the marshaling, manipulation and communication of information represented in a wide variety of formats. (Glasgow, et. al. 1995 presents a recent collection of influential articles in this area.) For example, an electrical engineering team may use circuit diagrams, state machine diagrams, timing diagrams and ladder logic diagrams, together with algebraic or natural language specifications of the desired input/output behavior, to produce a design that meets the client's needs (Harel 1988; Johnson, Barwise & Allwein 1996). Until recently, theoretical accounts of reasoning have limited themselves to homogeneous linguistic reasoning, that is, reasoning in which all information is represented using sentences of some natural or formal language. This limitation presents a major obstacle to the application of insights about the structure of reasoning to real-world, collaborative problem solving (Sloman 1975; Funt 1980; Barwise & Etchemendy 1991; Lewis 1991; Barwise 1993).
- **Heterogeneity of rationale.** The existing formal model of a natural deduction proof is far too constrained to countenance the kinds of reasoning involved in design and

[1] There is of course a long-standing controversy among psychologists about the viability of natural deduction as a model of human reasoning. This controversy is irrelevant to the present work, however, since both sides of the debate (e.g. Johnson-Laird & Byrne 1991, Rips 1994) have focused on the aspects of natural deduction that we set aside here, namely the inference rules associated with the so-called logical constants. In contrast, Keith Stenning and his collaborators have argued that the real correspondence between natural deduction and human performance can be found at the "case" or "subproof" level (Stenning 1998) that we have found useful in modeling heterogeneous reasoning.

practical problem solving. In particular, the model is too strict to countenance forms of justification based on matters of cost, efficiency, safety, style, aesthetic judgment, probabilistic considerations, and the like, justifications that arise at almost every turn in real-world problem solving (Mitchell 1990, e.g.). Some of these forms of justification are computationally tractable, and are therefore candidates for computer-aided verification, while others depend essentially on human judgment. Extending the natural deduction model to incorporate heterogeneous rationales is particularly important for collaborative reasoning, where a consensus must be achieved in spite of competing justifications (say aesthetic versus structural) of divergent decisions.

- **Heterogeneity of goals.** Historically, logic has focused almost exclusively on a single type of reasoning in which the goal is to show that a conclusion is a necessary consequence of some given information. In real-world problem solving, by contrast, the goals of a particular reasoning task can be quite varied. (Barwise & Etchemendy 1994, e.g.) Planning problems, for example, typically admit of a wide variety of solutions that meet the primary specifications, and selection among competing solutions is made on the basis of subsidiary, comparative criteria. The primary goal is thus to find *any* solution that satisfies certain requirements, not a unique solution entailed by those requirements, and secondary goals guide our choice among solutions to the primary goal. Research in logic has historically not addressed reasoning with such complex goal structures.

These three forms of heterogeneity, while conceptually distinct, are closely related. For example, if the goal of a reasoning task is to maximize a design solution along a certain metric, then the rationale for various design decisions will naturally involve that metric. Similarly the satisfaction of a design goal, say the final design of a building, will typically require a great variety of representations, both graphical and textual. We use the term *heterogeneous reasoning* to refer to reasoning that displays one or more of these three interrelated forms of heterogeneity.

These forms of heterogeneity in reasoning are all exhibited by *Hyperproof* (Barwise & Etchemendy, 1994), a text/software package aimed at instruction in analytical reasoning and problem solving. *Hyperproof* was the first natural deduction proof environment to allow heterogeneous forms of representation, in this case sentences of the first-order predicate calculus together with iconic representations of worlds populated by blocks located on a two-dimensional grid. The program allows students to construct heterogeneous proofs and automatically verifies the validity of their reasoning. *Hyperproof* broke with tradition by integrating heterogeneous representations, rationales, and goals into the natural deduction framework.

Figure 1 shows a simple proof constructed using *Hyperproof* illustrating all three forms of heterogeneity. The problem solved begins with heterogeneous information in the form of the diagram at the top of the figure together with three textual premises expressing constraints on the situation depicted by the diagram. The diagram shows an arrangement of six blocks of unknown size and shape, and one block of unknown size that is known to be a dodecahedron. The sentences say that two objects are the same size if and only if they are in the same row, the same shape if and only if they are in the same column, and that the further forward they are, the smaller they are. The problem has five goals, of four distinct types, shown at the bottom of the figure. The proof breaks into two distinct and exhaustive cases, or subproofs, based on the second premise plus domain specific constraints governing the *Hyperproof* blocks worlds. The user has continued reasoning within each case using the other premises. The short proof establishes all five goals. The entire proof, including its goals, has been verified by the *Hyperproof* program. (A detailed description of this proof is available at http://www-vil.cs.indiana.edu/TARK/SampleProof.html.)

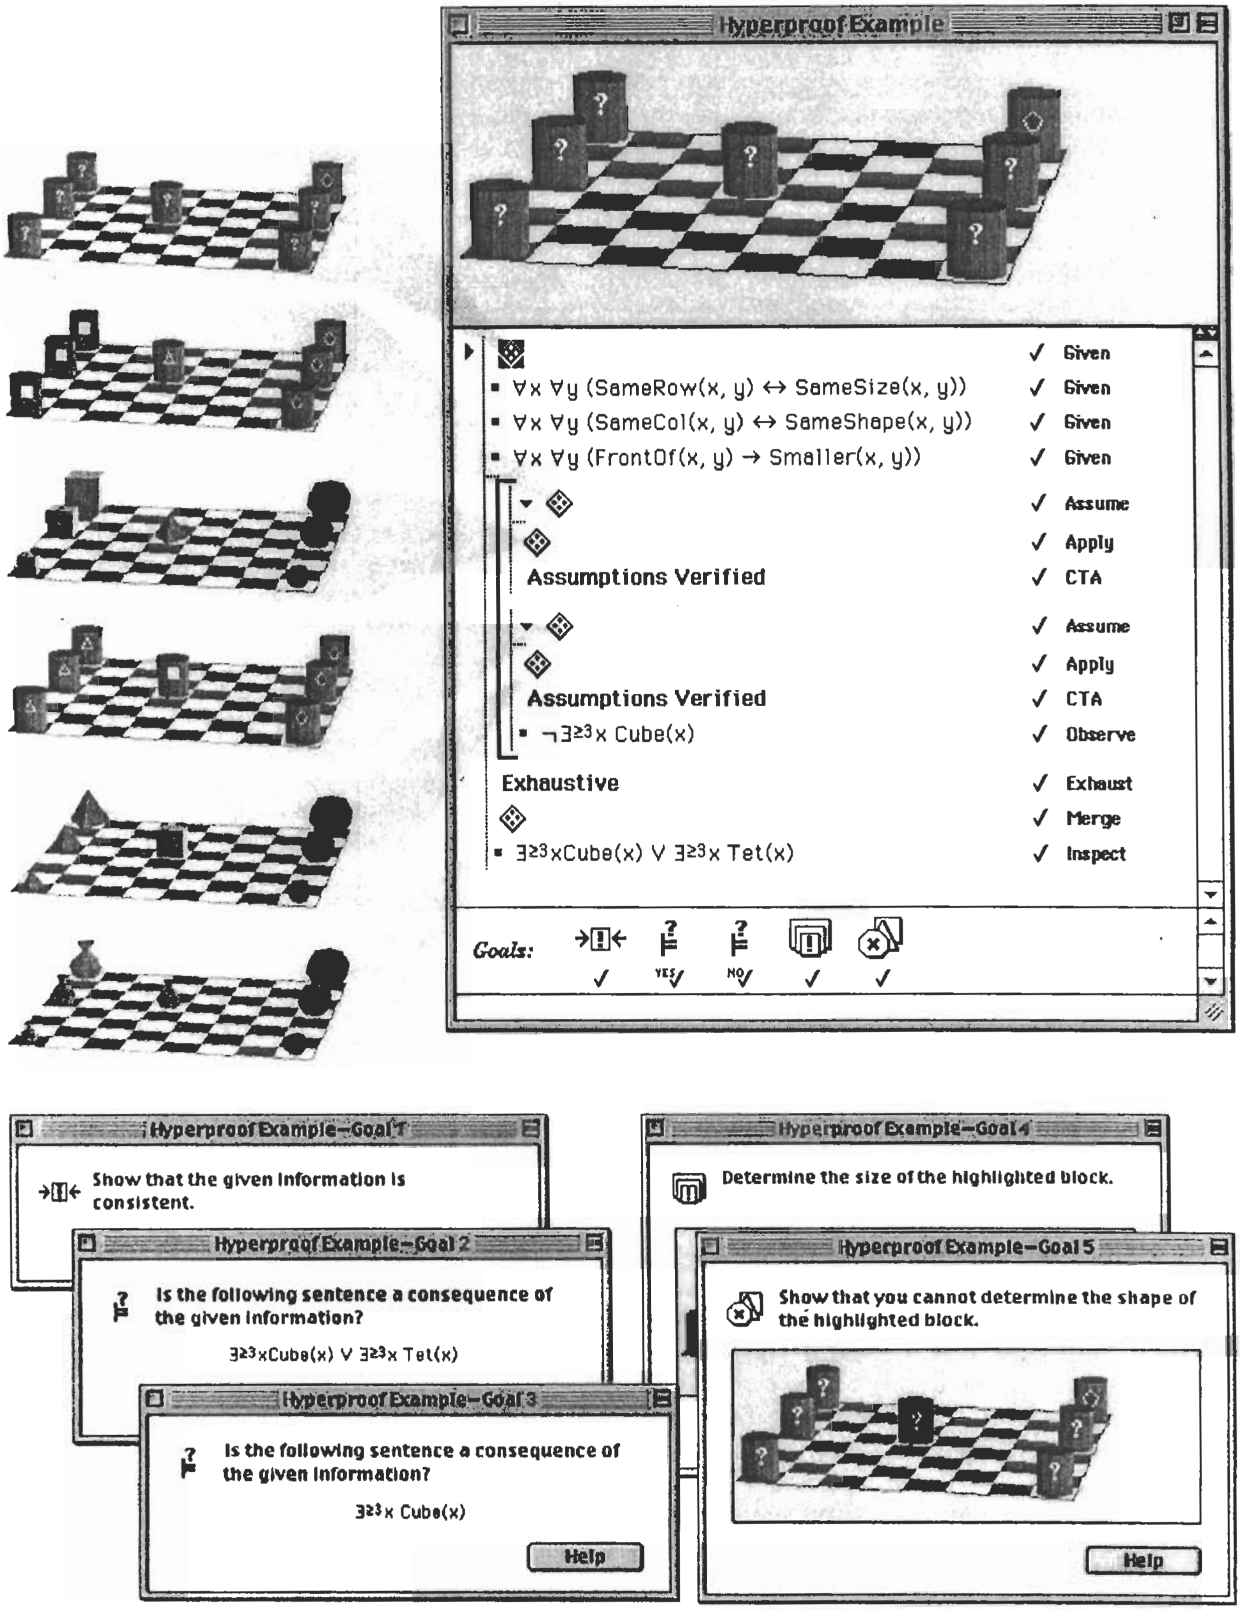
Hyperproof Example
∀x ∀y (SameRow(x, y) ↔ SameSize(x, y))
∀x ∀y (SameCol(x, y) ↔ SameShape(x, y))
∀x ∀y (FrontOf(x, y) → Smaller(x, y))
Given
Given
Given
Given
Assume
Apply
Assumptions Verified
CTA
Assume
Apply
Assumptions Verified
CTA
¬∃≥3x Cube(x)
Observe
Exhaustive
Exhaust
Merge
∃≥3xCube(x) ∨ ∃≥3x Tet(x)
Inspect
Goals:
YES
NO
Hyperproof Example–Goal 1
Show that the given information is consistent.
Hyperproof Example–Goal 2
Is the following sentence a consequence of the given information?
∃≥3xCube(x) ∨ ∃≥3x Tet(x)
Hyperproof Example–Goal 3
Is the following sentence a consequence of the given information?
∃≥3x Cube(x)
Help
Hyperproof Example–Goal 4
Determine the size of the highlighted block.
Hyperproof Example–Goal 5
Show that you cannot determine the shape of the highlighted block.
Help

Figure 1

Notice that the overall structure of this reasoning fits neatly into the familiar subproof style of reasoning familiar from natural deduction. Notice, too, though, that none of the rules cited in support of individual steps is one of the introduction or elimination rules associated with the logical constants. The rules used are all "formal" rules in that they can be checked by the computer, and they are sound rules, but they are not rules of the form familiar from sentence-based logical theory.

Our claim that natural deduction provides a useful framework in which to think about and build support tools for design and problem-solving often strikes those familiar with these formalisms as startling, due to the traditional emphasis on rules tied directly to syntactic characteristics of a formal language. As emphasized above, and illustrated in the example, the aspects of natural deduction that prove valuable in the general setting are its recursive mechanisms for managing the global informational dependencies within an extended, often complex, piece of reasoning. The approach provides a natural model of reasoning that respects the recursive structure of reasoning and the hierarchical nature of justification.

Openproof

Impressed by the naturalness of the reasoning permitted by *Hyperproof* and the effectiveness of the program as a pedagogical tool, we began developing a general theory of heterogeneous reasoning and exploring extensions of the basic architecture to more realistic problem domains. We mention in particular the dissertations of our graduate students Ruth Eberle (*Hyperproof*), Kathi Fisler (hardware circuits), Mark Greaves (geometry), Eric Hammer (higraphs, among others), Mun-Kew Leong (maps), Isabel Luengo (geometry), Sun-Joo Shin (Venn diagrams) and Atsushi Shimojima (general properties of diagrammatic representations), and Michael Wollowski (diagrams as used in planning). Some of this work by these individuals is sampled in the collection Allwein & Barwise 1996.

Our own work in this area resulted in a patent application (Etchemendy & Barwise 1998) for a general architecture for recording heterogeneous reasoning in the form of "proofs" (in our extended sense of the term). **Figures 2, 3**, and **4** illustrate three potential applications of this architecture. The purpose of these examples is to give a sense of the variety of reasoning to which the architecture may be applied. The examples themselves are all artificial and highly simplified for the sake of illustration.

Figure 2 shows a simple proof prepared by an architect designing an addition to an existing house (shown at step 1). The client wants to add a guest bedroom and bath; the architect's proposed solution is shown at step 3. This solution is the result of a structured reasoning process that is recorded in the proof. The proof is a record of the rationale for the decisions incorporated into the final design, and can be used to present the reasoning to colleagues and clients. (For a discussion of the relevance of *Tarski's World* and similar tools to the architectural design process, see Mitchell 1990.) This reasoning is heterogeneous in all three of the senses explained above, and illustrates an application of the architecture in a domain where one would not expect computer verification of many justifications.

Figure 3 also stresses the heterogeneity of the reasoning and justification process. The reasoning recorded concerns the assignment of offices, represented diagrammatically in the proof, based on a variety of constraints that are expressed sententially. The reasoning shows that there is only one possible assignment that satisfies the stated constraints. If there had been more than one solution, then secondary considerations could be brought to bear in selecting the optimal solution. In this case, verifying the primary justifications is computationally tractable, while verifying secondary considerations might well not be.

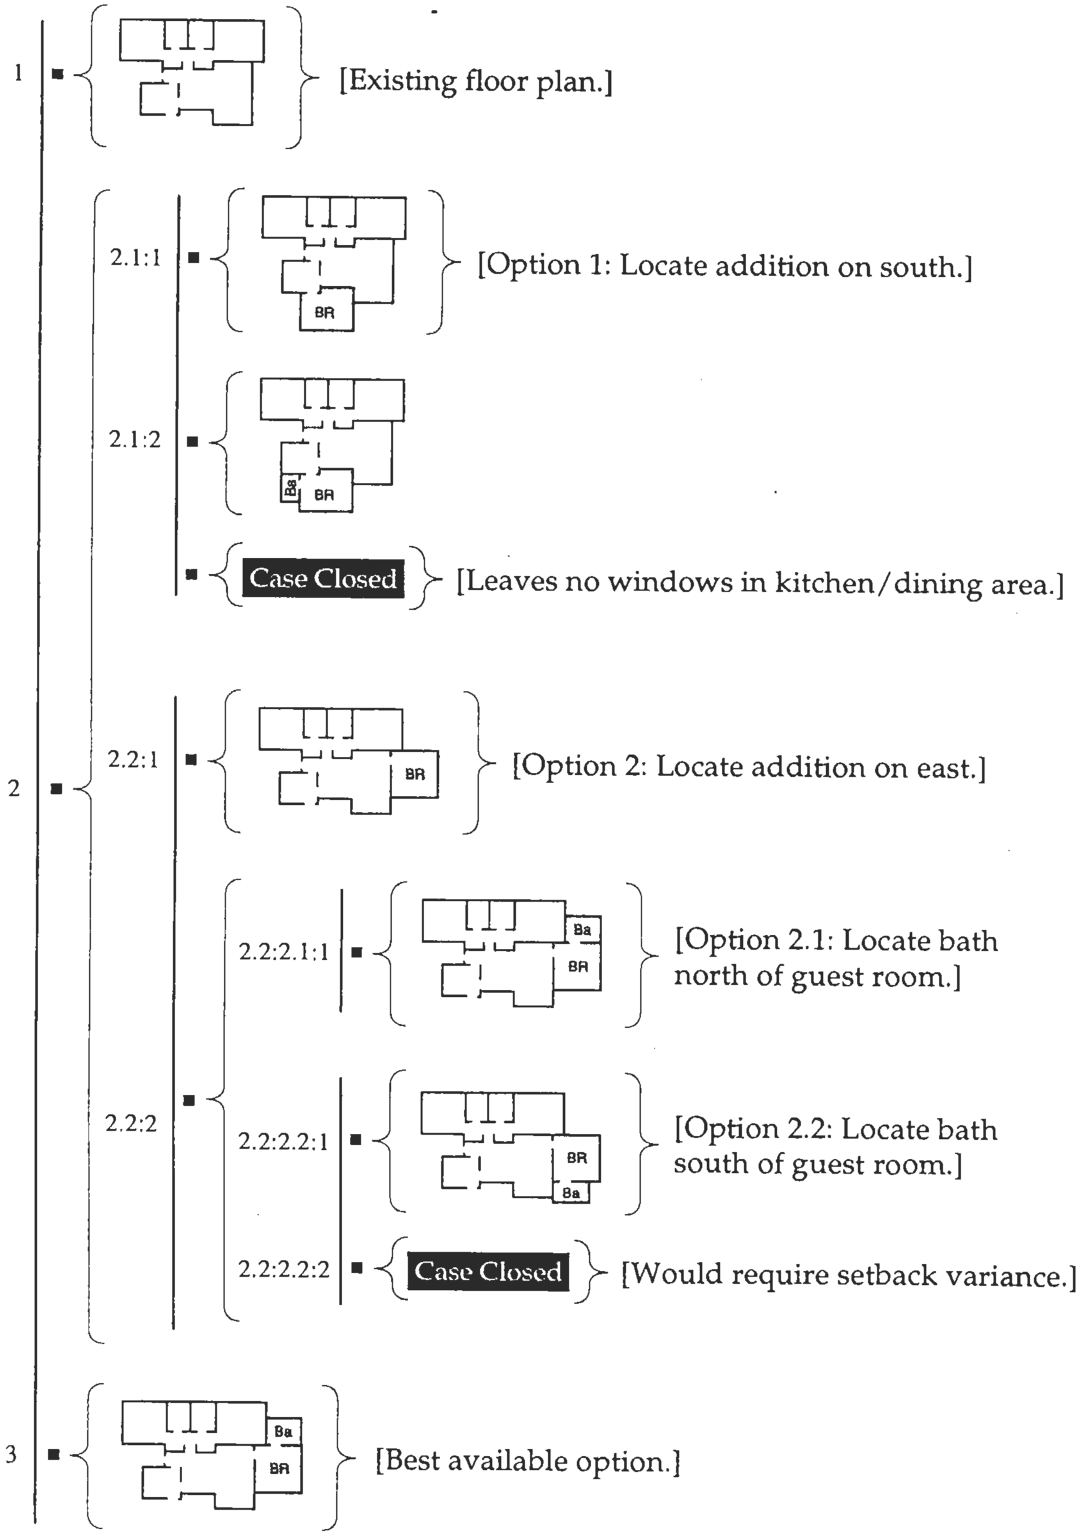
1
[Existing floor plan.]
2.1:1
BR
[Option 1: Locate addition on south.]
2.1:2
Ba
BR
Case Closed
[Leaves no windows in kitchen/dining area.]
2
2.2:1
BR
[Option 2: Locate addition on east.]
2.2:2
2.2:2.1:1
Ba
BR
[Option 2.1: Locate bath north of guest room.]
2.2:2.2:1
BR
Ba
[Option 2.2: Locate bath south of guest room.]
2.2:2.2:2
Case Closed
[Would require setback variance.]
3
Ba
BR
[Best available option.]

Figure 2

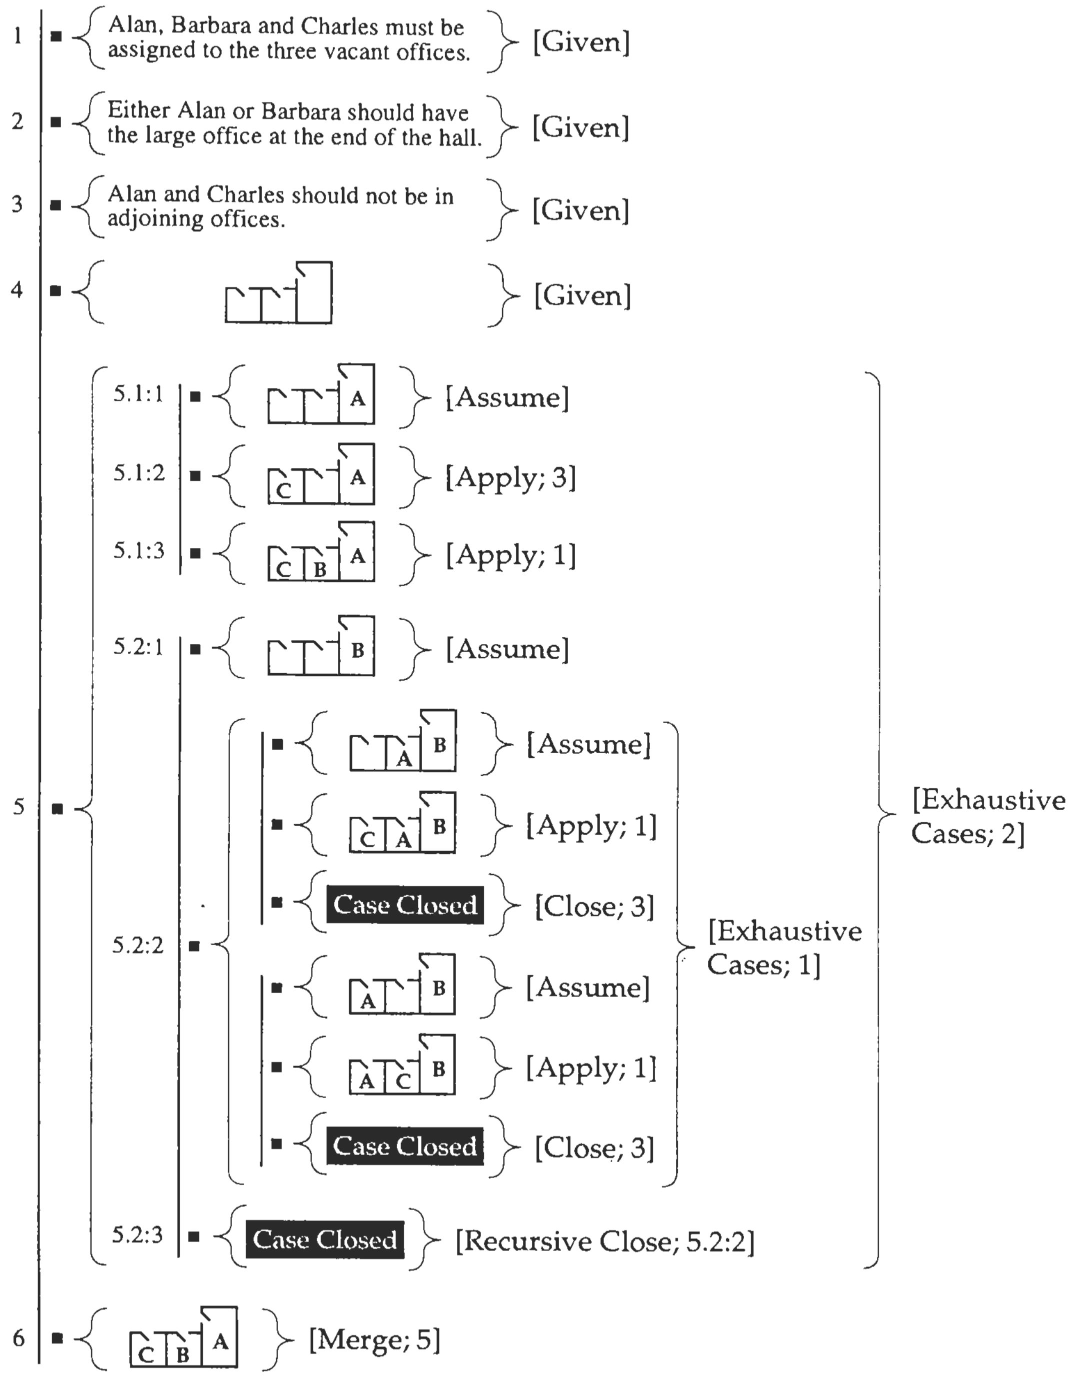
1 Alan, Barbara and Charles must be assigned to the three vacant offices. [Given]
2 Either Alan or Barbara should have the large office at the end of the hall. [Given]
3 Alan and Charles should not be in adjoining offices. [Given]
4 [Given]
5.1:1 A [Assume]
5.1:2 C A [Apply; 3]
5.1:3 C B A [Apply; 1]
5.2:1 B [Assume]
A B [Assume]
C A B [Apply; 1]
Case Closed [Close; 3]
5.2:2
A B [Assume]
A C B [Apply; 1]
Case Closed [Close; 3]
[Exhaustive Cases; 1]
5.2:3 Case Closed [Recursive Close; 5.2:2]
5
[Exhaustive Cases; 2]
6 C B A [Merge; 5]

Figure 3

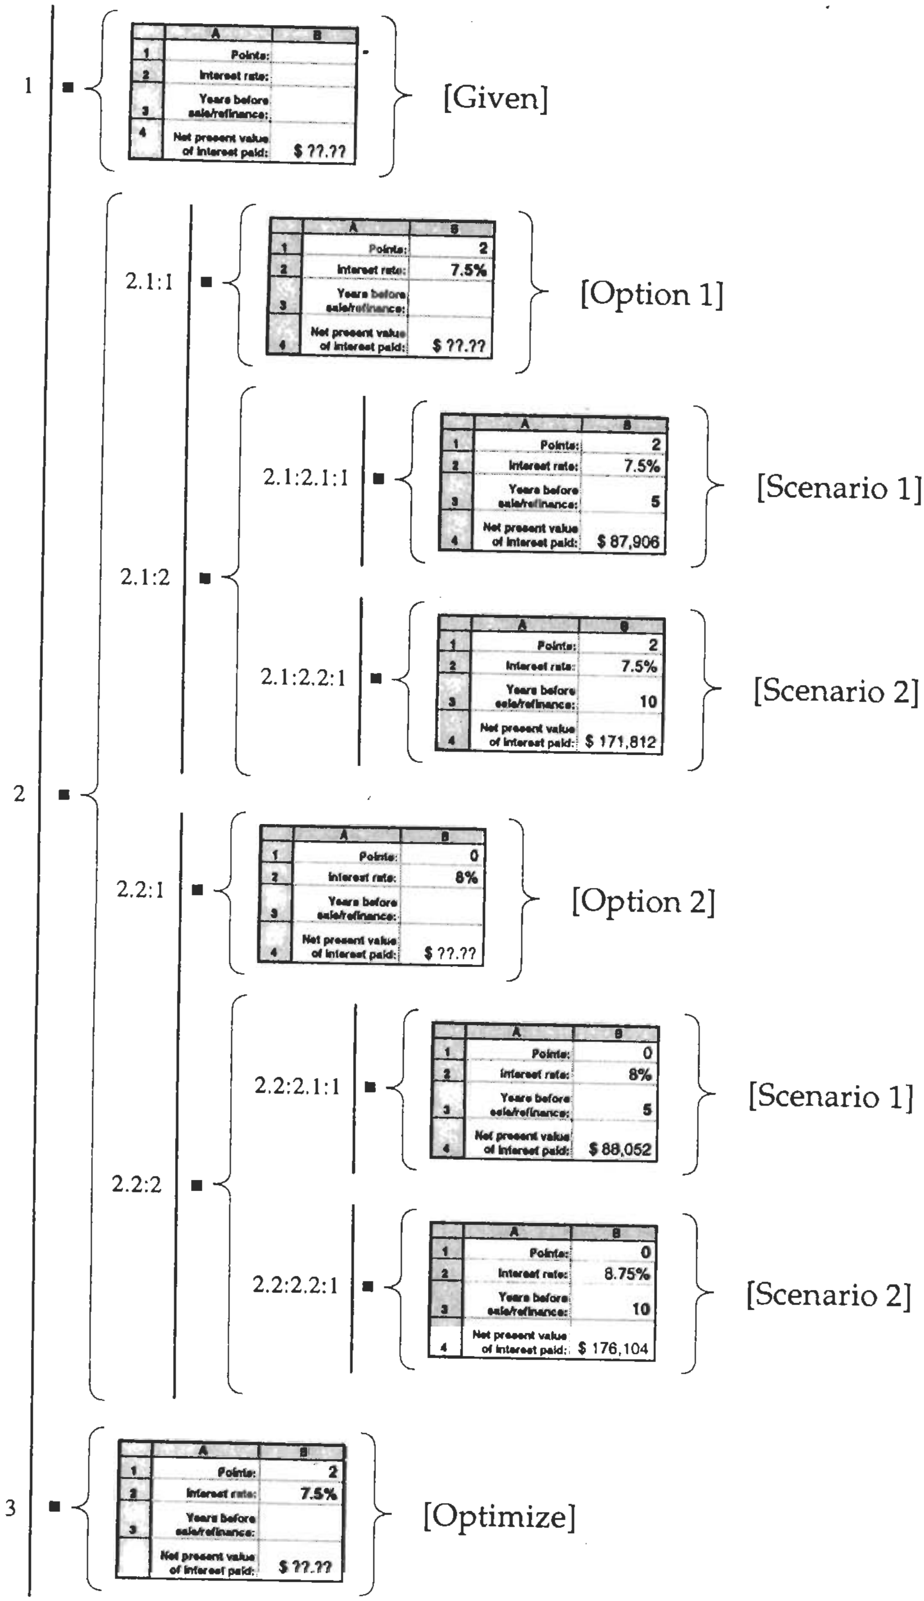

1
Points:
Interest rate:
Years before sale/refinance:
Net present value of interest paid: $??.??
[Given]
2
2.1:1
Points: 2
Interest rate: 7.5%
Years before sale/refinance:
Net present value of interest paid: $??.??
[Option 1]
2.1:2
2.1:2.1:1
Points: 2
Interest rate: 7.5%
Years before sale/refinance: 5
Net present value of interest paid: $ 87,906
[Scenario 1]
2.1:2.2:1
Points: 2
Interest rate: 7.5%
Years before sale/refinance: 10
Net present value of interest paid: $ 171,812
[Scenario 2]
2.2:1
Points: 0
Interest rate: 8%
Years before sale/refinance:
Net present value of interest paid: $??.??
[Option 2]
2.2:2
2.2:2.1:1
Points: 0
Interest rate: 8%
Years before sale/refinance: 5
Net present value of interest paid: $ 88,052
[Scenario 1]
2.2:2.2:1
Points: 0
Interest rate: 8.75%
Years before sale/refinance: 10
Net present value of interest paid: $ 176,104
[Scenario 2]
3
Points: 2
Interest rate: 7.5%
Years before sale/refinance:
Net present value of interest paid: $??.??
[Optimize]

Figure 4

Figure 4 illustrates an application of the architecture to the very different domain of financial planning. In this example, the user is deciding between a mortgage at 7.5% interest, with two points, and one at 8% interest, with no points. The user envisages two salient scenarios: one in which he is transferred out of town in five years; one in which he keeps the house for ten years. The reasoning shows that under either scenario, the optimal choice is the former mortgage. This example is particularly suggestive in the use of what we dub a "metric promotion rule" at step 3. In this example, the user has a dominant strategy, but the architecture can support reasoning that incorporates probabilities and allows rationales involving expected utility.

Although the reasoning illustrated in these figures has been simplified for the sake of illustration, it is important to understand that the architecture supports very complex reasoning and design tasks. Most crucially a single "proof" can maintain large numbers of non-textual representations of multiple types as they evolve through the reasoning or design process, and record their informational dependencies. For example, a single step in a proof may contain circuit diagrams, timing diagrams, and so forth, along with justifications for changes made at that step to one or more of these representations. A subproof can trace the ramifications of a design decision until it is sufficiently elaborated (perhaps by means of recursively embedded subproofs) to assess its satisfaction of project goals or to allow comparison with alternative decisions.

We believe that the architecture will prove valuable in large collaborative projects in at least four ways. First, it structures the project evolution in a minimally intrusive manner, one that will seem natural to users. Second, it captures design intent and rationale, both for immediate communication to other project members and for creation of a permanent, easily comprehensible project memory. Third, it allows multiple users to work simultaneously on different parts of a project, to add rationale from different disciplines, or to satisfy different goals. Finally, by supporting multiple representations within a single structure, it allows project members to use different representation types while maintaining crucial links between the information conveyed in those representations as the design evolves.

Our long-term goal is to refine and implement this architecture. We believe it has the potential to affect a wide variety of engineering and scientific practices by improving the processes through which designs and problem solutions are created, assessed, communicated, and recorded in collaborative settings. It is potentially applicable in any domain in which reasoning employs one or more forms of non-textual representations, which includes virtually every scientific and engineering discipline. It will also have a significant impact on the theoretical understanding of reasoning processes and on pedagogical practices in the teaching and learning of analytical problem solving skills.

As a first instance of this architecture we are working with Gerard Allwein and David Barker-Plummer on the development of *Openproof*. This consists of a domain-independent natural deduction proof manager plus communication protocols for modular representation editors and drivers designed to support reasoning appropriate to specific application domains. These protocols permit domain-specific and representation-specific inference rationales, and flexible goal specification and satisfaction mechanisms. Computer-aided verification of justifications and goal satisfaction can be incorporated into modules where feasible.

Barker-Plummer, Etchemendy, and Renate Fruchter have begun investigating the incorporation of heterogeneous reasoning technology into Fruchter's project management software. We plan to use the insights gained in this collaboration to build *Openproof* components for a prototype architectural CAD system.

We also plan to develop new courseware tentatively named *The Sand Reckoner*. This will replace *Hyperproof* as a tool for teaching and learning heterogeneous reasoning skills. Where *Hyperproof* has only one non-textual mode of representation, *The Sand Reckoner* will have several. The current plan is to develop *Openproof* modules for blocks world diagrams, Venn and Euler diagrams (Shin 1994, Hammer & Shin 1996), position diagrams (Barwise & Etchemendy 1989, 1991), Boolean truth tables, binary decision diagrams, and free-form sketches. In addition, *Sand Reckoner* will allow problems employing both symbolic formulas and natural language sentences.

In addition, we are working with researchers at the Boeing Company to develop a set of *Openproof* modules to help programmers author reliable intelligent agent applications. These modules will employ several types of domain specific graphical representations to enable system development teams to reason about conversations that take place in agent-based systems, and so to improve the reliability of such systems.

Other related research

While we have focused on our work and that of our research team, this work has not taken place in a vacuum. Nontextual representations have been used in collaborative human reasoning since one of our ancestors first picked up a stick and sketched a plan for a hunt. Until recently it has been studied by only a few logicians, including Euler, Venn, and Peirce, but by many psychologists, including Arnheim, Koslyn, Johnson-Laird, and Stenning. Interest in diagrammatic reasoning appeared in artificial intelligence in the 70's and 80's (Sloman 1971, 1975; Funt 1980, e.g.) but it was only in the 1990's that the full-fledged study of nontextual reasoning started to come into its own. Recent cognitive and computational work in this area is well represented in the collection (Glasgow, Narayanan & Chandrasekaran 1995). Of particular relevance to this project are the papers (Harel 1988), (Myers & Konolige 1995), (Stenning & Inder 1995), and (Wang, Lee & Zeevat 1993).

We called for the study the logical properties of nontextual and heterogeneous reasoning in Barwise & Etchemendy 1989, 1991. Bailin and Barker-Plummer have investigated the use of diagrams in mathematical practice (Barker-Plummer & Bailin 1997). The basic observation in this work is that while mathematicians usually deny any formal role for diagrams within proofs, many presentations of mathematical proofs include diagrams as an aid to human intuition. Bailin and Barker-Plummer have developed a theory that suggests that mathematical diagrams often suggest cut formulae, or lemmas, that constrain the search for the proof. They have interrogated this theory by implementing a completely automatic reasoning system (&/GROVER) that embodies the theory, and by comparing the ability of that system to complete difficult proofs in mathematics both when presented with a supporting diagram and without such a diagram. &/GROVER can automatically complete proofs of some quite difficult theorems (by automated reasoning standards) when presented with the diagram, and not when the diagram is omitted. These results demonstrate that useful guidance can be drawn from a diagram by automated reasoning systems, and suggests that something analogous is happening in the human case.

We have worked with Keith Stenning and his colleagues in studying the learning effect of *Hyperproof*. Stenning predicted that students' learning benefits from the *Hyperproof* diagrams would be determined by their facility with the graphical abstraction "tricks" which must be used in order to pose insightful heterogeneous reasoning problems. This prediction was born out by analysis of the logs of activity from the first class taught using *Hyperproof*, as compared to a control class taught using sentential natural deduction (Stenning, Cox & Oberlander 1995). Further analysis of this data has demonstrated that the determining factor in students' learning from the heterogeneous environment is their strategic insight into the issue of when to translate information from sentences to diagrams, and when from diagrams to sentences (Oberlander, Cox, Tobin, Stenning & Monaghan 1996). This result has been generalized to students' spontaneous diagram drawing in solving analytical reasoning problems (Cox, Stenning & Oberlander 1995); and to other domains of reasoning where it is possible to study on-line tutoring dialogues (Monaghan & Stenning (submitted)). Out of this line of work emerges computational foundations for a theory of the differences in human reasoning styles that underlie individuals' responses to heterogeneous reasoning and communication systems.

Conclusions

Students often leave a logic course feeling that logic is a puerile activity unrelated to everyday life. Viewed as the science of rational justification, however, it ought to be central to every field of science and engineering as well as law, business and medicine. We feel that expanding the conception of logic from the study of homogeneous deductive inference based on the meaning of a handful of logical constants, to that of heterogeneous reasoning as defined here has the potential to transform both the subject and the workplace. Our ambition is to contribute to this transformation. We hope that our theory of heterogeneous reasoning, the *Openproof* framework, and the prototype systems mentioned above will advance the theory and practice of collaborative, heterogeneous reasoning, and pave the way for the application of these methods to other application domains. While the problems involved in large scale collaborative design and problem solving processes are enormous, given the pervasiveness, importance and cost of this activity, we believe that even a modest contribution to its efficiency and reliability has a large potential payoff.

Acknowledgments

This paper evolved from a grant proposal to the NSF. We acknowledge the contributions to the ideas here by other investigators on the team, including Gerard Allwein, Dave Barker-Plummer, Kathi Fisler, Renate Fruchter, Mark Greaves, and Keith Stenning.

Hyperproof was conceived by us and implemented by a team led by Allwein and Greaves. The software has been extended and modified by Allwein and Barker-Plummer, and is currently in use at institutions around the world. (The software is described at http://csli-www.stanford.edu /hp/index.html.) The positive educational impact of the software has been in Cox & Oberlander 1995; Cox and Stenning & Oberlander 1995. Allwein and Barker-Plummer have are collaborators in the design and implementation of *Openproof*. Mark Greaves is leading the Boeing part of the collaboration mentioned above.

References

Allwein, G. and J. Barwise (1996). *Logical Reasoning with Diagrams*. Studies in Logic and Computation. New York: Oxford University Press, 270 pages.

Bailin, S.C. (1997). "Software development as knowledge creation." International Journal of Applied Software Technology, vol 3, no 1, March, 1997.

Barker-Plummer, D. and S. C. Bailin (1997). "The Role of Diagrams in Mathematical Proofs," in *Machine Graphics and Vision*, vol 6, no 1, 25-56.

Barwise, J. (1993). "Heterogeneous reasoning," in *Conceptual Graphs and Knowledge Representation*, ed. by G. Mineau, B. Moulin, and J. F. Sowa, Lecture Notes on Artificial Intelligence, Cambridge: MIT Press, 64-74.

Barwise, J. and J. Etchemendy (1986/1993). *Turing's World: An Introduction to Computability*. First published by Santa Barbara: Academic Courseware Exchange (1986). Stanford: CSLI and Cambridge: Cambridge University Press (1993), 123+ix.

Barwise, J. and J. Etchemendy (1987/1991). *Tarski's World.* First published by Santa Barbara: Academic Courseware Exchange (1987). Stanford: CSLI and Cambridge: Cambridge University Press, first edition (1991) with Macintosh software, 111+xv; second edition (1993) with Macintosh software, 116+xviii; second edition (1993) with IBM software, 122+xviii.

Barwise, J. and J. Etchemendy (1989). "Information, infons, and inference," in *Situation Theory and its Applications I*, ed. by Cooper, Mukai, and Perry, Stanford: CSLI and Cambridge: Cambridge University Press, 33-78.

Barwise, J. and J. Etchemendy (1991). "Visual Information and Valid Reasoning," in *Visualization in Mathematics*, ed. Walter Zimmerman and Steve Cunningham, MAA. 9-24.

Barwise, J. and J. Etchemendy (1992). "Hyperproof: Logical Reasoning with Diagrams," in *Proceedings of the 1992 AAAI Spring Symposium on Diagrammatic Reasoning,* Stanford: AAAI, 1992, 80-84. Reprinted in *Reasoning with Diagrammatic Representations,* Menlo Park: AAAI Press, 1994.

Barwise, J. and J. Etchemendy (1994). *Hyperproof.* Program by Gerard Allwein, Mark Greaves, and Mike Lenz. Stanford: CSLI and Cambridge: Cambridge University Press, 255+xvii.

Barwise, J. and J. Etchemendy (1995). "Heterogeneous Logic," in (Glasgow et. al. 1995). 211-234.

Barwise, J. and J. Etchemendy (1998). "Computers, Visualization, and the Nature of Reasoning," in *The Digital Phoenix: How Computers are Changing Philosophy.* Ed. T. W. Bynum and James H. Moor. Blackwell. (Also available at http://www.phil.indiana.edu/~barwise/CV&NR.pdf.)

Chandrasekaran, B., A. Goel, and Y. Iwasaki (1993). "Functional representation as design rationale." *IEEE Computer*, vol. 26, no. 1, 48-56.

Conklin, Jeff and Begeman, Michael (1989). "IBIS: A Tool for All Reasons," *JASIS: Journal American Society for Information Science*, 40:3, 200-213.

Cox, R., K. Stenning, and J. Oberlander (1995). "The effect of graphical and sentential logic teaching on spontaneous external representation." *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society*, 2(4), 56-75.

Etchemendy, J. and J. Barwise (1998). "Computational Architecture for Reasoning Involving Extensible Graphical Representations," pending U.S. patent application, filed 05/98.

Fisler, K. (1996). "Exploring the potential of diagrams in guiding hardware reasoning," in Allwein and Barwise (1996), 225-256.

Fisler, K. (1996b). *A Unified Approach to Hardware Verification Through a Heterogeneous Logic of Design Diagrams*. PhD Dissertation. Indiana University Department of Computer Science, August 1996.

Fitch, F. B. (1952). *Symbolic Logic: An introduction*. Luvain, 238 pages.

Fruchter, R. and Krawinkler, H. (1995). "A/E/C Teamwork," Proc. *Second ASCE Congress of Computing in Civil Engineering*, Atlanta, June 1995, 441-448.

Fruchter, R. (1996). "Conceptual, Collaborative Building Design Through Shared Graphics," *IEEE Expert Intelligent Systems*, June 1996 Vol 11 no. 3, 33-41.

Fruchter, R., Reiner, K., Leifer, L., and Toye, G. (1998). "VisionManager: A Computer Environment for Design Evolution Capture," accepted for publication in *CERA: Concurrent Engineering: Research and Application Journal*, Vol 6, no. 1, 1998.

Funt, B. (1980). "Problem-solving with Diagrammatic Representations," *Artificial Intelligence* 13, 1980, reprinted in Glasgow, et al. 1995, 33-68.

Glasgow, J., H. Narayanan, B. Chandrasekaran (1995). *Diagrammatic Reasoning: Cognitive and Computational Perspectives*, AAAI Press and MIT Press, 780 pages.

Gentzen, G. (1935/1969). "Investigations into logical deduction," In M.E. Szabo (Ed. and trans.) *The collected papers of Gerhard Gentzen.*, Amsterdam: North-Holland.

Hammer, E. (1995). *Logic and Visual Information*, Studies in Logic, Language and Information, Stanford: CSLI and Cambridge: Cambridge University Press.

Hammer, E and S-J Shin (1996). "Euler and the role of Visualization in Logic." *Logic, Logic, Language and Computation I*, J. Seligman and D Westerstahl, 271-286.

Harel. D. (1988). "On Visual Formalisms," Communications of the ACM 13 (5) (May 1988), 514-530, reprinted in Glasgow, et al. (1995), 235-272.

Johnson-Laird, P. N. and R. M. J. Byrne (1991). *Deduction*. Hove, UK: Lawrence Erlbaum Associates.

Kosslyn, S. M. (1994). *Image and brain : the resolution of the imagery debate*. Cambridge, Mass: MIT Press.

Mitchell. W. J. (1990). *The Logic of Architecture*. Cambridge: MIT Press.

Monaghan, P. and K. Stenning (submitted) "Effects of representational modality and thinking style on learning to solve reasoning problems," *20th annual meeting of the Cognitive Science Society of America.*

Myers and Konolige (1995). "Reasoning with analogical representations," in Glasgow, et al. (1995), 273-302.

Lewis, M. (1991). "Visualization and situations," in *Situation Theory and its Applications, vol* 2, ed. by Barwise, Gawron, Plotkin and Tutiya, Stanford: CSLI and Cambridge: Cambridge University Press, 553-580.

Oberlander J., R. Cox, R. Tobin, K. Stenning and P. Monaghan (1996). "Individual differences in proof structures following multimodal logic teaching." Garrison W. Cottrell (ed.) *Proceedings of the 18th Annual Conference of the Cognitive Science Society of America*, UCSD, La Holla, Ca. July 1996. 201-206.

Parnas, D. and P. Clements (1986). "A rational design process: how and why to fake it." IEEE Transactions on Software Engineering, SE-12(2), February 1986.

Prawitz, D. (1965). *Natural Deduction. A proof-theoretical study.* Almqvist and Wiksell: Stockholm, 113 pages.

Prawitz, D. (1973). "Towards a foundation of a general proof theory," *Proceedings of the 1971 International Congress of Logic, Methodology, and Philosophy of Science.*, Bucharest, 225-250.

Rips, L. J. (1994). *The Psychology of Proof.*, Cambridge: MIT Press.

Shin, S.-J. (1994). *The Logical Status of Diagrams*, Cambridge: Cambridge University Press.

Sloman, A. (1971). "Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence." In Proceedings of the Second International Joint conference on AI, London. Morgan Kaufmann.

Sloman, A. (1975). "Afterthoughts on analogical representation," in *Readings in Knowledge Representation.*, ed. R. J. Brachman and H. J. Levesque. Morgan Kaufmann.

Stenning, K. and Inder (1995). "Applying semantic concepts to analysing media and modalities," in Glasgow, et al. (1995), 303-338.

Stenning, K. and J. Oberlander (1995). "A cognitive theory of graphical and linguistic reasoning: logic and implementation," Cognitive *Science*, 19, 97-140.

Stenning, K., R. Cox and J. Oberlander (1995). "Contrasting the cognitive effects of graphical and sentential logic teaching: reasoning, representation and individual differences," *Language and Cognitive Processes*, 10, 333-354.

Stenning, K. and O. Lemon (in press). "Aligning logical and psychological perspectives on diagrammatic reasoning." *AI Review*.

Stenning, K. and P. Yule (1997). "Image and language in human reasoning: a syllogistic illustration." *Cognitive Psychology*, 34 , 109-159.

Stenning, K. (1998). "Aligning logical and psychological contributions to the understanding of human reasoning." *Kognitionswissenschaft*, 7(1).

Wang, D., Lee, J. and Zeewat (1995). "Reasoning with diagrammatic representations," in Glasgow at al (1995), 339-393.