

On hypermedia-based argumentation decision support systems ¹

Gary H. Hua ^{a,2}, Steven O. Kimbrough ^{b,*}

^a *Research and Development Department, Reed Technology and Information Services, One Progress Drive, Horsham, PA 19044, USA*

^b *University of Pennsylvania, The Wharton School, Steinberg Hall–Dietrich Hall, Suite 1300, Philadelphia, PA 19104-6366, USA*

Abstract

This paper presents and discusses a logical apparatus which may be used to support machine-based inferencing and automatic creation of hypertext links in what we call hypermedia-based argumentation decision support systems (HADSS). This logical approach has important advantages over other sorts of argument representation, found in the current literature. We present and discuss a prototype implementation in the context of three examples. We also present an exploratory experiment indicating that graph-based logical representations can materially help people make better inferences. © 1998 Elsevier Science B.V. All rights reserved.

Keywords: Argumentation; Decision support systems; Defeasible reasoning; Hypermedia; Hypertext; Logic

1. Introduction

There is a growing consensus, supported by a growing literature, to the effect that principles of hypermedia (including hypertext and multimedia), when applied in argumentation decision support systems (DSSs), will add enormous value. ³ The theory here is that hypermedia systems add value: by facilitating information filtering (e.g., with graphical representations of arguments), by assisting in information retrieval (e.g., at the click of a mouse), and by facilitating vivid and forceful presentation of information (e.g., with animation and business graphics). The consensus—with which we are in accord—is

widely shared and is expressed in a broadly-based literature. Argumentation concepts have been articulated (e.g., see Refs. [22,30,31,36]), ^{4,5} and systems have been built (e.g., Refs. [1,9–11,22,42,43,49]).

This consensus—on the need for, and practicality of, hypermedia-based argumentation DSS (HADSS)—is supported by two sorts of evidence. First, existing systems, while mainly prototypes, have generally been convincing to those who have seen them and worked with them. People, at least many people, have found the idea of hypermedia-based argumentation DSS very attractive, once they have seen examples of such systems. This evidence, however, has circulated more in the oral tradition than in the literature. Second, there is a small but growing body

* Corresponding author. Tel.: +1-215-898-5133; e-mail: kimbroughwharton.upenn.edu.

¹ File: HADSS-DSSSI-19970908, HADSS-DSSSI-199771019.

² Tel.: +1-215-682-4289; e-mail: ghua@reedtech.com.

³ An argumentation DSS is a DSS aimed at helping people to construct, evaluate, and communicate arguments [36].

⁴ See also H.V.D. Parunak, Ordering the Information Graph, in Ref. [3], pp. 299–325.

⁵ See also E.B. Swanson, Business Value as Justificatory Argument, in Ref. [58], pp. 121–138.

of empirical research showing that people, unaided by supporting technology, perform rather poorly at constructing, evaluating, and communicating arguments (e.g., see Refs. [2,37] and Ref. [24], especially Chapter 2). This work can be seen as doing for logic and argumentation what Kahneman et al. [27] and Tversky and Kahneman [61,62] have done for utility theory: demonstrating that, unaided by appropriate technology, people's performances fall considerably short of what the best available normative theory prescribes.⁶

Our aims in this paper are to raise three central questions about HADSSs (there are others), and to advance (admittedly partial) answers to these questions: (1) How should arguments (including rebuttals, counter-arguments, rebuttals of rebuttals, and so on) be represented? See Section 2; (2) What features can and what should be provided by way of machine support for creating, evaluating, comparing, and exploring arguments? See Section 3. (3) How useful and practicable is the HADSS idea? See Section 4, where we discuss examples suggesting that a positive answer to this question can, with further research, be made; see Section 5, where we present an experiment that indicates that with proper graphical representation people make fewer (deductive) inferential mistakes?

2. Arguments and argumentation theory

Reasoning has been described as “mental activity that consists of transforming given information (called the set of premises) in order to reach conclusions” (Ref. [18], p. 333). In the jargon of logic, the conclusion is a particular statement, and the premises are also statements (linguistic expressions that can be either true or false). Together, they are called an argument. It is the main business of logic to study arguments. As reasoning is the transformation of

premises to reach a conclusion, then we may characterize a reason as an account of how particular premises were transformed to produce a particular conclusion. Equivalently, then, we may say that it is the main business of logic to study reasons.

An argument, thus, from the point of view of logic, consists of zero or more premises, and a conclusion. Both premises and conclusion are statements, possibly complex statements that are composed from other statements. In a successful argument, the premises can be taken as good reasons to believe the conclusion, because it can be established that if the premises are true, then the conclusion must be true. Arguments with this property are said to be valid. If the premises of a valid argument are also true, then the argument is said to be sound. Logic, to repeat, is the study of arguments, of what makes some better than others, of what makes an argument successful or not, of what constitutes validity, and so on. The actual truth or falsehood of premises is usually taken not to be part of logic, but part of natural science, broadly conceived. Similarly, the actual causal processes we undergo in transforming premises to conclusion is usually seen as part of cognitive psychology. Consequently, the theory of arguments and argumentation is a very broad one, and has (at least) a logical aspect and a psychological aspect. We shall touch on both aspects in what follows.

There are two traditions in the study of logic: formal logic and informal (or semi-formal) logic. Formal (or symbolic or mathematical) logic is the study of logic in which the arguments under consideration are symbolized into a well-defined language that can be manipulated according to a strict set of rules and independently of the meanings (or intended interpretation) of the sentences in the argument symbolized. Logic was first studied from a formal point of view by Aristotle and has been studied, and developed, more or less continually since his time. Precision and rigor have been the motivations for the formal study of logic, and a rich variety of fully formal logical calculi have been developed and examined. A happy consequence of this formalization is that the resulting calculi can be—and have been—implemented in software.

Formal logic, then, can be thought of as providing a computational and a normative theory of argument.

⁶ This research tradition is beginning to intersect with argumentation research. Behavioral decision theorists—and behavioral game theorists—are increasingly looking to behavioral argumentation theory to explain observed departures from what decision theory—and game theory—predict (H. Kunreuther, personal communication, 1992; J.D. Laing, personal communication, 1992).

This is just what is needed, or so it would seem, for computerized argumentation support systems. There are problems, however. First, as might be expected of any effort at formalization, there are problems of limitation in scope (we call this the scope problem). Although the range of formal logical languages is impressive and growing, it still falls well short of the requirements of everyday, practical deliberation. Things as simple as newspaper articles (particularly editorials and op ed pieces) severely stretch the natural expressive capabilities of the standard logics, which do not do well at representing such concepts as time and obligation [35]. Much worse, in our view, is that standard logics do poorly at handling defeasible reasoning, reasoning that involves inferences that may be defeated by new information (See, e.g., Refs. [7,19,20,38,45–49,52,53,55] for discussion of defeasible reasoning.⁷ See Refs. [23,29,32–34] for a technical introduction to the methods of logic graphs and of sweeping presumptions, the logic we have developed for argumentation DSSs.) Since nearly all practical inferences are defeasible (some examples, including a newspaper example, are presented in Section 4), this is indeed a worrisome shortcoming.

Formal logics have a second serious problem, if they are to serve as bases for computerized argumentation support systems. Logical notation is arcane (we call this the friendliness problem). Without extensive training in logic, it is simply not on to expect users to have meaningful understanding of formalized expressions.

These problems have long been recognized and have led to the study of logic from a non-formal, or semi-formal, point of view. Such studies have been conducted since the time of Plato. There are broadly two relevant literatures. First, there is an extensive literature that goes by the name of informal logic. There are two primary journals in this field, *Informal Logic*, and *Argumentation*, as well as two professional societies, the Association for Informal Logic and Critical Thinking, and the International Society for the Study of Argumentation. In addition, quite a

number of textbooks have appeared devoted to the subject (e.g., see Refs. [12,14,15,26,28,50,54,59,60] as well as Refs. [63–66]). The second body of literature is much smaller. It consists of reports of experiences with implementations of argumentation systems. Examples include Refs. [1,4,11,21,22,36,39–44,56,57].⁴

It is not to our purpose here to review this literature in any depth. What matters most, for present purposes, is that both literatures are essentially in agreement that arguments can (and, for pragmatic reasons, should) be represented as directed graphs, in which nodes stand for statements and arcs stand for certain ‘semantic’ relations between statements.^{8,9} For example, Hashim’s (Ref. [22], p. 238) list of semantic relations for Issue-Based Information Systems (IBISs) includes ‘supports’, ‘challenges’, ‘responds to’, and ‘objects to’, but there is no broad agreement on these labels.

We may compare formal and informal (or semi-formal) approaches to argument representation in computerized support systems. Formal logics are beset with the scope and friendliness problems. Informal (or semi-formal) approaches are agreed on resolving the friendliness problem by representing arguments as directed graphs, although different informal approaches will differ on how these graphs are to be constructed. Informal approaches may be thought of as dealing with the scope problem by treating the semi-formal, graph-based representation as an approximate model of the underlying argument. Users viewing the representation are then free to make transformations and manipulations that reflect intuitive concepts of, say, time, obligation, and defeasibility. The graph representation is used much as a rough map: not itself a calculation device, it helps by prompting the user to think things through.

The informal approaches, thus, have much to be said in their favor. What they lack is a principled

⁷ See also M.L. Ginsberg, Introduction, in Ref. [20]; D. Nute, Defeasible Logic and the Frame Problem, in Ref. [38], pp. 1–21; and D. Nute, Defeasible Logic, in Ref. [17], pp. 353–395.

⁸ This statement itself needs qualification, especially with regard to the problem of representing quantification, but for present purposes this detail can safely be ignored. We refer the reader to the treatment in Ref. [34].

⁹ We note that in an early and intriguing article, Swanson also represented arguments as graphs (⁵). It is a quite natural thing to do.

theory of inference. Using the IBIS framework (but the point generalizes), if statement A supports B, which supports C, if D challenges C, and if both A and D are presumed true, then what are we to believe about C? There is no theory to help us, nor is there any calculation that might be offered in this regard. Addressing the scope problem by going informal is a sword that cuts two ways.

Ideally, we would like graph-based argumentation representations with an underlying proof theory that addresses the problem of scope. This is not a world in which ideals are often realized. We submit, however, that logic graphs and sweeping presumptions (discussed in Refs. [23,33,34,29]) are fair contenders. Fair enough, at least, to be worth exploring. The main purpose of the remainder of this paper is to report upon our initial investigations of using logic graphs (for deductive inference) and sweeping presumptions (an extension of logic graphs, for defeasible reasoning) as a graph-based representation scheme for argumentation DSSs.

3. Machine assistance

The problems of scope and friendliness hardly exhaust the challenges in designing and implementing a HADSS. Paramount among the remaining challenges is the treatment of hypermedia.

A HADSS, as we conceive it, would have two broad categories of hypermedia links: links between elements of the argumentation network (internal links), and links between elements in the network and elements outside the network (external links). For example, at a given node, the presence of an incoming arc might be exploited to determine that another node (the node originating the incoming arc) provides support for the statement at the node we are at, i.e., the given node. This would be an internal link. An external link at a node might be used to display corroborating documents, videos, and so on for the assertion at the node.

As we have noted (indeed, emphasized), in logic graphs and sweeping presumptions arguments are representable as networks. This fact can be exploited to counter a well-known problem with hypermedia: the high cost of building hyperdocuments due to manual node creation and linking [5,6,9].

Any HADSS should facilitate interactive, ad hoc construction and modification of arguments, and provide hypertext-style access to documents relevant to particular arguments. These aims would strongly suggest that a gIBIS-style (graphical IBIS, see above) interface [11] could be usefully adapted for browsing through argumentation representations. In gIBIS-style systems, the user is presented with a graph in which the nodes and arcs are given special meaning, pertaining to the subject at hand. For example, in an argumentation system (as noted earlier) nodes might represent issues, positions, or arguments, and arcs might represent such ‘semantic’ notions as ‘supports’, ‘rebutals’, and ‘suggests’.

We agree that such a facility (a gIBIS-style interface for a HADSS) would very likely be useful. Merely having such an interface, however, leaves unresolved two important problems.¹⁰

The first problem is that of building a system that can automatically construct hypertext links. Specifically, documents will be relevant to particular nodes and arcs. A user should be able to direct the system to produce, say, all the documents relevant to the node representing objections to the ‘pay or play’ proposal for health care financing. Additionally, the HADSS should, insofar as is possible, find these documents automatically, that is without a person having explicitly to specify that a particular document is relevant. In previous works, we have addressed this problem successfully, but in a different context [5,6,36]. We believe, however, that the techniques developed and applied in that work can be useful in the present context. Because the nodes (and the arcs) are meaningful (they represent propositions or relations among propositions) and formally represented, they can be exploited to create external links automatically. For example, if a node asserts that P, then documents that discuss P are possibly relevant to that node. Since P is a formal expression, text generation techniques can be used to generate key words, which may then be used to search the database

¹⁰ These two problems, in our view, also apply to other existing systems that support argumentation in some way, e.g., as those in Refs. [21,39,57] and other systems cited above. The fundamental problem is lack of a theory of inference for the argument representation. See our discussion above, in Section 2.

of documents. Further, the location of node P in the network can be exploited to focus the search for relevant documents. For example, if Q implies P and Q is known (or presumed to be true), then it may be prudent to focus the automated search on documents pertaining to both P and Q. Generalizations of these points are easy to imagine. It remains to test them on real databases and apply them to real problems.

The second problem (with relying only on gIBIS-style interfaces) has to do with a limitation in existing gIBIS-style systems, which use the IBIS framework for representing arguments (and other things, such as discussions) graphically. For the purpose of modeling arguments, the IBIS framework is but one of many (see references above), and it has some important shortcomings. For present purposes, the most salient limitation of the IBIS framework is that it is not a logic and does not support logical inference. At best, it would be difficult to support machine-based inference for an IBIS representation.

Thus, as an alternative to IBIS, we are proposing to use logic graphs and the method of sweeping presumptions to represent arguments in HADSS implementations.¹¹ A HADSS could be given a number of operations that help with decision support by exploiting the network representation (and underlying theory, sweeping presumptions) to produce links automatically. These operations include procedures to answer the following questions. Answers could be assembled dynamically and displayed as new hyper-text nodes.

(1) Is the argument valid? This operation checks if an argument is constructed according to the argument formation rules (syntax validity checking). The operation also checks if there are any conflicts among the subarguments and if all the assumptions are supported by the evidence (semantics validity checking). In the case of a finding of invalidity, a recommendation can be proposed.

(2) Are there counter-arguments to a given argument? This operation generates a set of potential counter-arguments to an argument from the evidence

and reasonable assumptions. A counter-argument is any rebutting defeater or undercutting defeater to the argument (See, e.g., Refs. [51] for definitions of these terms).

(3) Is a counter-argument effective? This operation compares the relative strengths between the conflicting arguments, and recommends some methods to enforce an argument. For example, the strength of an argument can be increased by increasing support to a critical premise or by increasing the strength on the defeasible rules in the argument or by introducing the defeaters to the counter-argument.

(4) What if we changed a given assumption? ‘What-if’ analysis is the process of making changes to the evidence data, the assumption data or the strengths of defeasible inference rules and observing the impact to an argument and its counter-arguments.

(5) What would need to be the case for this argument (or counter-argument) to be effective (goal-seeking analysis)? Goal-seeking analysis is the process of providing a reasoning chain for an argument which supports a desired conclusion. This operation identifies the critical evidence, assumptions, as well as the defeasible inference in the arguments. The argument can be weakened or invalidated by weakening the strengths of the defeasible inferences or by removing the critical evidence or assumptions. Alternatively, the irrelevant information for an argument can also be identified.

We have developed a prototype implementation of a HADSS, and we call it HADSS. This prototype is based on sweeping presumptions, and is able to draw theoretically-correct logical inferences. Thus, it directly addresses the validity question (1, above) for the arguments it represents. Our implementation also directly addresses the effectiveness of the counter-argument question (see 3, above). Counter-arguments are represented as rebutting and/or undercutting defeaters and the program totals up the strengths (as determined by sweeping presumptions) of the reasons pro and con. The what-if question (4, above) is directly addressed by our HADSS. Users may interactively set statements (represented as nodes) to be true, false, or not known to be either. In virtue of having the what-if capability, our HADSS has modest support for the goal-seeking question (5, above). Users may seek goals by iterating what-if questions. Finally, there is some computational support for the

¹¹ We note that there are other graph-based logic representations that could serve as alternatives to our particular methods. Notable in this regard are the works of Nute and Erk [48] and Nute et al. [49].

counter arguments question (2, above). HADSS can list the rebuttals and defeaters (see Ref. [23]) for a given conclusion. We note, however, that at present the listings become hard to follow if the underlying arguments are complex. Users must often rely on visual inspection of the argument graph. We plan to address this problem in a future version.

4. Implementation

We will describe our HADSS implementation in the context of three examples.

4.1. Toulmin's Anne

Toulmin (Ref. [59], p. 126) presents a simple, but interesting and commonsensical, example of defeasible reasoning: (1) Anne is one of Jack's sisters; (2) All Jack's sisters have previously been observed to have red hair; (3) So, presumably, Anne now has red hair, unless Anne has dyed her hair, gone white, lost her hair, etc.

Toulmin has his own diagrammatic approach to representing arguments, but for the sake of conserving space we shall not further discuss his approach here. Under our approach, this argument might be graphed as in Fig. 1.¹²

Briefly, our implementation works as follows. HADSS allows the argument graph to be constructed interactively. Once the network is constructed, the user may declare certain assumptions and perform the operations for answering the questions listed in Section 3. In the present example, one declares the assumption that Anne is one of Jack's sisters by clicking on the $S(x)$ node and responding to a dialog box. If there are no further assumptions, the user may then click on the conclusion node, $NR(x)$, and ask whether it can be inferred. The system then produces a report regarding the assertion at this node. Here, HADSS would report that $PPNR(x)$ can

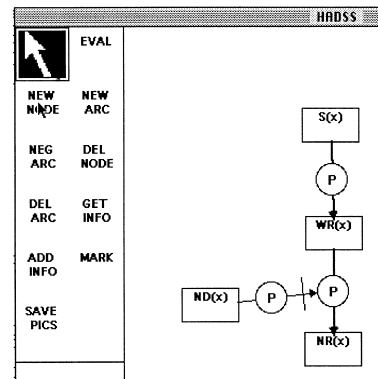


Fig. 1. The HADSS representation of Toulmin's Anne: $S(x)$, x is one of Jack's sisters; $R(x)$, x has red hair; $D(x)$, x has dyed hair, $P\phi$, presumably, ϕ ; $W\phi$, it was true at least once that ϕ ; $N\phi$, it is now true that ϕ .

be derived ('Presumably, presumably it is now true that $R(x)$ ') with a valid argument which supports $NR(x)$, and that nothing else can be derived concerning $NR(x)$.¹³ Suppose the user adds the assumption that Anne now has dyed hair (this is done as above, by clicking on the appropriate node). If then we click on the conclusion node, HADSS reports (see Fig. 2) that $PPNR(x)$ can be derived (as before), but also reports this argument is defeated by an undercutting defeater. Consequently, nothing can be concluded regarding the color of Anne's hair. HADSS provides users with a graphical interface for constructing argument networks and for making initial assumptions, so the users can easily perform what-if analysis by adding or deleting links and changing the assumptions. Further, since the strengths of relevant arguments are explicitly displayed, the users may analyze how changes on strengths of competing arguments may effect the conclusion.

4.2. Presidential polls

Our second example comes from an op ed piece, which discusses why (on June 27, 1992) the various

¹² In this example, we are ignoring issues about scope of quantification on logical variables. This is not a fundamental problem for our approach. The issue is easily handled, but requires more logical sophistication than is warranted in the present paper. See Ref. [34] for details.

¹³ In sweeping presumptions, iterations of 'presumably' operators weaken the statement. Thus, $PP\phi$ indicates less credence for ϕ than does $P\phi$.

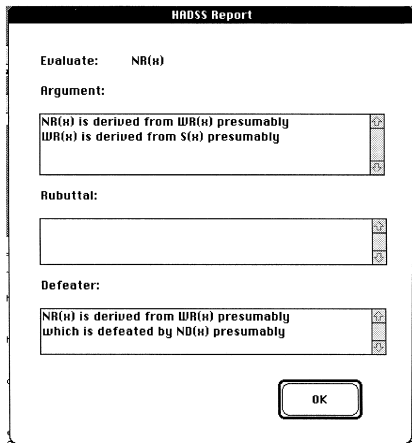


Fig. 2. The HADSS report on Toulmin's Anne.

presidential polls were in so much disagreement [16]. In essence, the thesis of the article is that the polls are in disagreement (with each other) because the voters' judgments are volatile. The author of the article, Kathleen Frankovic, assumes that if voters are volatile (VV) then presumably the polls will be in mutual disagreement (PD). She provides evidence (more on this shortly) that in fact the polls are in disagreement. Most of the article (which is quite short) is given over to providing evidence regarding the fact that different questions are asked by different polling organizations (DQDA) and that these lead to the different answers we see. At the end of her article, Frankovic briefly asserts that the polls also show that voters are distressed and confused (VDC). Her argument may be summarized with the argument network in Fig. 3, which is a screen dump from the HADSS representation of the network.




Fig. 3. The HADSS representation of presidential polls argument: DQDA, different questioning by different polls produces different answers; PD, the polls disagree; VV, the voters are volatile; VDC, the voters are distressed and confused.

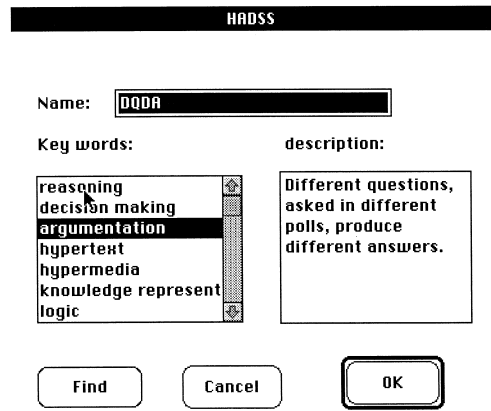


Fig. 4. The HADSS representation of information attached to node DQDA in the presidential polls argument.

The features described above with regard to Toulmin's Anne (Section 4.1) apply here as well. The argument in the network is a simplified and abstracted version of that presented in the article. For example, Frankovic presents specific claims about how The New York Times/CBS News poll was taken, about the Time–CNN poll, and so on. This information is not represented in the argumentation network per se, but it is available to the HADSS user. Sticking to our present example, the polling information is relevant to the argument because it supports the assertion at the DQDA node. A user may gain access to this information by clicking on the DQDA node and asking for more information (see Fig. 4).

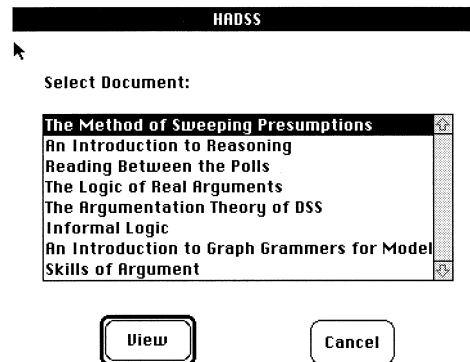


Fig. 5. The HADSS representation of documents associated with the keyword argumentation in the presidential polls argument.

In the current implementation more information is obtained inferentially (and dynamically, at run time) by using key word indices to extract a list of relevant documents from a document description database. That list is then presented to the user (see Fig. 5), who may then select detailed information on a particular document. In Fig. 6, we see that in Frankovic’s article polling methods are discussed for the New York Times/CBS polls, the Time–CNN polls, and the Gallup Organization polls. Although our implementation does not currently do this, it would be a simple matter to allow the user to view a particular document at this point.

4.3. PNB’s money access center (MAC) decision

Money Access Center is one of the country’s largest ATM networks and has been dominant in the Philadelphia region, where its owner, Philadelphia National Bank (PNB), is located. MAC is by all accounts a high-quality, profitable system with an excellent business record and good prospects for continuing success. Surprisingly, however, MAC is a shared network—many competing banks participate in it—owned by a single bank (PNB), which competes with most of the network’s customers. Further, MAC was launched in the presence of a competing ATM network—Girard Bank’s George—that had high visibility and extensive coverage in the Philadelphia area.

The story of how MAC was conceived, launched and brought to success is a fascinating one, rich in implications for IS and strategy. This story, moreover, has been ably recounted by Clemons [8]. Our

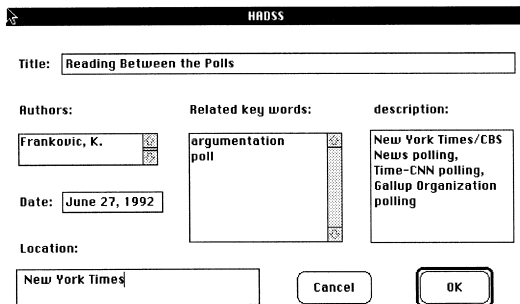


Fig. 6. The HADSS representation of information about a certain document linked to the presidential polls argument.




Fig. 7. Stage 1 of the PNB MAC decision: the problem. Because Girard’s George ATM network is successful (GGIS), if PNB does not get a large ATM network (¬LATM), then presumably PNB will have a serious loss in its retail banking business (SLRBB).

purpose here is not to repeat the story, but to extract from it something of the argumentation structure presented to PNB’s decision makers prior to the decision to launch MAC as a shared network. Space limitations confine us to a discussion of the argumentation network, but the HADSS functionality should be apparent.

According to Clemons [8],¹⁴ strategy formulation by PNB regarding MAC may be reconstructed as a three-stage process. In stage 1, PNB recognized a problem. Girard Bank had successfully launched the George ATM network, and PNB was concerned that its retail customers would move in droves to Girard, unless PNB developed a similarly large ATM (LATM) network. The situation is summarized in Fig. 7.

In stage 2 of the policy formulation process, options were identified for consideration. The first thing assumed was that PNB would have a LATM network available to it. In terms of our argumentation network, this is the denial of ¬LATM. Similarly, ¬SLRBB (no significant loss of retail banking business) is added as a node, since that is a main objective of the policy formulation exercise. See Fig. 8.

A critical assumption at stage 2 was that if PNB is to have an LATM, then (presumably) it would have to build its own private network (BONET). Given this, there would seem to be two main options: build a network with about as many ATM machines as Girard George’s (1TO1) or (xor node in Fig. 8) build a much smaller network with about the same ratio of customers to ATM machines as Girard’s George (PCENT). The problem with these options is that neither leads to averting a significant loss of retail

¹⁴ E.K. Clemons, personal communication, July 1992.

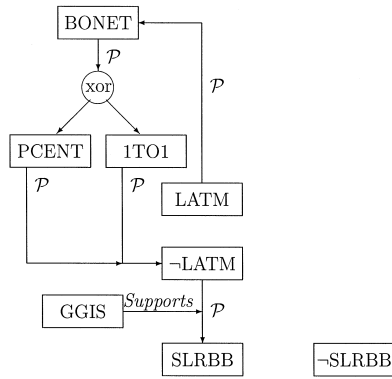


Fig. 8. Stage 2 of the PNB MAC decision. BONET, PNB builds its own ATM network; ITO1, PNB matches the number of ATMs fielded by Girard's George; PCENT, PNB fields an ATM network, matching the machines to customer ratio of George.

banking business (\neg SLRBB), as is clear in Fig. 8. Both options fail to produce an LATM network for PNB, PCENT by design and ITO1 because it is simply infeasible due to cost.¹⁵

In terms of sweeping presumptions and the argumentation network in Fig. 8, if we assume that PNB is to have an LATM network, then we are led to conclude that either PPPSLRBB (if ITO1 is chosen) or, again, PPPSLRBB (if PCENT is chosen). These are horns of an unhappy dilemma, and they in fact correspond to what common sense reports (and PNB's management saw) about this situation. Is there a way out? The diagram is helpful. Consider the assumptions that are being made at stage 2. Clearly, nodes GGIS and LATM are assumed, and clearly they really are not at all controversial in this context. There are other assumptions, however. Each arc represents an assumption. For example, Fig. 8 assumes (defeasibly) that if BONET, then there are exactly two options available, PCENT and ITO1. This is certainly false, but not in any way that matters (Argumentation is a form of modeling, too. Like other forms, successful practice requires tasteful

¹⁵ Since node \neg LATM is also reached, it might appear that we have fallen into contradiction. We have not, due to the use of defeasible reasoning and sweeping presumptions. There is, e.g., no contradiction between 'Presumably, Q' and 'Presumably, \neg Q.' Here, and often in general, we are simply recognizing that what evidence there is conflicting. For details, see Ref. [34].

and judicious embracing of the false in favor of the feasible.)

What Clemons reports (personal communication) is that the LATM–BONET arc was, in effect, recognized by PNB as meaningfully defeasible. The arc is assuming that PNB will have, as Girard did, a single-bank network. Once this conceptual shift was made, it was fairly clear what the options were.

There were two alternatives to reduce cost while meeting PNB's very demanding requirements for retail service delivery: to reduce cost by shared development of a multiple-owner, consortium ATM network, or to reduce final cost to PNB by developing the type of network PNB desired and aggressively marketing it to other Philadelphia area banks. The second alternative, marketing of services, was seen to have better 'up-side' potential. Moreover, the consortium development was seen as impractical at the time: most banks that were able to commit resources were developing or planning to develop proprietary networks, and consortium development appears very difficult to coordinate. No more detailed formal analysis was done, and the decision to launch and market a single-owner network was made in 1978 (Ref. [8], emphasis added).

This reasoning is modeled in Fig. 9. Given the model, it is clear why MAC wins over CNET, assuming we are using sweeping presumptions. We

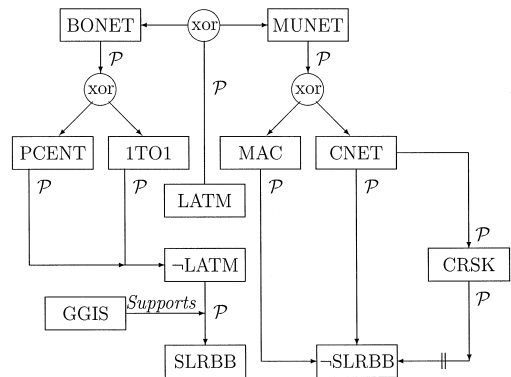


Fig. 9. Stage 3 of the PNB MAC decision. MUNET, PNB participates in a multiuser ATM network; MAC, PNB builds the MAC ATM shared network; CNET, PNB promotes a consortium ATM network; CRSK, the consortium network risks coordination difficulties. Note: || indicates denial; here, 'presumably not.'

conclude (with sweeping presumptions) that $PPP \rightarrow SLRBB$ (with MAC) or $PPP \rightarrow SLRBB$ but undermined with $PPPP \rightarrow SLRBB$ (with CNET). Thus, the weight of the argument favors MAC, and with sweeping presumptions we choose it, as (in effect) did PNB's management.

5. An experimental study

Argumentation DSSs are aimed at helping people construct, evaluate, and communicate arguments. Do they work? How can they best be designed and what might we expect from them? These are large, important, and difficult questions, questions that only a long stream of yet-to-be-done research can answer. Existing prototypes, including the work reported above, are mostly only prototypes. They are but the beginning of the exploration of an interesting idea. Our purpose in this section is to extend the conceptual/prototyping results (ours and others) by reporting on an experiment in which subjects were asked to make logical inferences either with or without aid of the sorts of argument representation-directed graphs-used in argumentation DSSs.

Substantial recent research has shown that people often do not reason in accord with the rules of formal logic. By way of explanation, Johnson-Laird presents a different point of view by introducing a theory, called 'mental models,' to describe how people solve deductive problems. Johnson-Laird (in Ref. [24], p. 98) observes that "... human reasoners appear to retain a superficial representation of the propositions expressed by the premises—one that is close to their linguistic form—but from the errors they make, they appear to make inferences by manipulating mental models rather than by deploying rules of inference on these superficial representations". Johnson-Laird shows that mental models can explain many aspects of human reasoning [24].

Despite the differences in the descriptions and models for reasoning processes, there is a general consensus that people should follow what the normative theory prescribes for deductive reasoning. One interesting observation is that people are normally willing to make corrections when they realize that they have violated the logical inference rules. It is not our purpose here to review and study the differ-

ent theories of reasoning. Our focus is on developing tools for supporting people in reasoning processes. As discussed above, many features in argumentation support systems, including graph-based representation of arguments, inference support mechanisms, and information (document) retrieval techniques, can support people in reasoning processes. In this section, we present an experiment to demonstrate that people can improve their deductive reasoning by using appropriate supporting tools. Because this is an initial experiment, and because of an extensive related literature, we have focused on deductive (rather than, say, defeasible) reasoning.

5.1. Experimental method

In the experiment, the subjects were presented with a set of deductive reasoning problems, and they were asked to derive conclusions from a set of given statements. The statements are conditional English sentences, which are similar to those Johnson-Laird and Byrne used in their experimental studies (see Refs. [24,25]). Previous studies (by Johnson-Laird, Byrne and others) have shown that people make systematic errors when they solve similar problems based on their intuitive judgments. In our experiment, we provide one group of subjects with a special tool, called the logic graph method [13,29]. As we discussed above, the logic graph method plays an important role in our logical framework. We developed the theory of argumentation support systems, as well as our prototype system, based on the ideas of the logic graph method and its extension, the method of sweeping presumptions. This experiment is designed to assess the usefulness of the logic graph method in deductive reasoning.

5.2. Subjects and procedures

The subjects were 57 University of Pennsylvania undergraduate students. None of the subjects had previously received any formal training in logic. Subjects were assigned randomly to two groups of approximately equal size: a control group (group (A), $N = 28$) and an experimental group (group (B), $N = 29$). The subjects in both groups were presented a set of problems. Each problem consisted of a set of English statements with a list of possible conclu-

sions. The subjects were asked to choose the correct conclusion(s) for each set of given statements. Further, in addition to the problem statements in English, the subjects in group (B) were presented with a logic graph representation of the problem, which they could use to help them in reasoning. Prior to the experiment, the subjects in each group were given 10 min to read training materials. To encourage the subjects to try their best to solve the problems, we offered a payment of US\$3.00 for each correct answer, for two randomly chosen subjects.

5.3. Training materials

5.3.1. Control group (group A)

The training material for group (A) consisted of examples and non-technical explanations on how to draw a conclusion from given statements. The following is a part of the training material. The complete training material is presented in Ref. [23].

Example:

Statement 1: If there is a circle then there is a triangle.

Statement 2: If there is a triangle then there is a square.

Statement 3: There is a circle.

Circle all the correct conclusions.

(a) There is not a circle.

(b) There is not a square.

(c) There is a triangle.

(d) There is not a triangle.

(e) None of the above is correct.

(Conclusion: (c))

Statement 1 is a conditional statement, in which, ‘there is a circle’ is the precondition, and ‘there is a triangle’ is the conclusion. Similarly, statement 2 is also a conditional statement, in which, ‘there is a triangle’ is the precondition, and ‘there is a square’ is the conclusion. In general, when the precondition of a statement is satisfied, its conclusion must follow. Therefore, when making inferences with conditional statements, we always check their preconditions first. As in this example, the given fact, ‘there is a circle’ (statement 3), satisfies the precondition of statement 1. We can draw the conclusion (c) ‘there is a triangle’. Further, since this conclusion, ‘there is a triangle’, satisfies the precondition of statement 2,

we can draw another conclusion ‘there is a square’ (which is not in the list of the choices).

5.3.2. Experimental group (group B)

Similarly, subjects in group (B) read training material consisting of examples, the logic graph representation of conditional statements, and non-technical explanations on how to draw a conclusion from the given statements with the logic graph method. The following is a part of the training material. The complete training material is presented in Ref. [23].

The logic graph method is a tool for representing statements and supporting users in making inferences. The representations of logical statements in the logic graph method are rather simple. We illustrate the logic graph method with three examples. Please read these examples very carefully.

Example:

Statement 1: If there is a circle then there is a triangle.

Statement 2: If there is a triangle then there is a square.

Statement 3: There is a circle.

Circle all the correct conclusions.

(a) There is not a circle.

(b) There is not a square.

(c) There is a triangle.

(d) There is not a triangle.

(e) None of the above is correct.

(Conclusion: (c))

Fig. 10 illustrates how to use the logic graph method to represent and solve the problem in the example. Step 1 depicts the logic graph representation of statement 1 and statement 2. The node with label ‘C’ represents ‘there is a circle’, and the node with label ‘T’ represents the statement ‘there is a triangle’. The ‘if-then’ relationship between ‘C’ and ‘T’ is represented by the arrow from node ‘C’ to node ‘T’. (a) in step 1 depicts the representation of statement 1—‘if there is a circle then there is a triangle’. Every conditional statement has a contrapositive equivalent statement. The contrapositive equivalent statement is just a different way to state the original statement. Sometimes, we find this equivalent statement is convenient to use in making inferences. The contrapositive equivalent statement of statement 1 is ‘if there is not a triangle then there

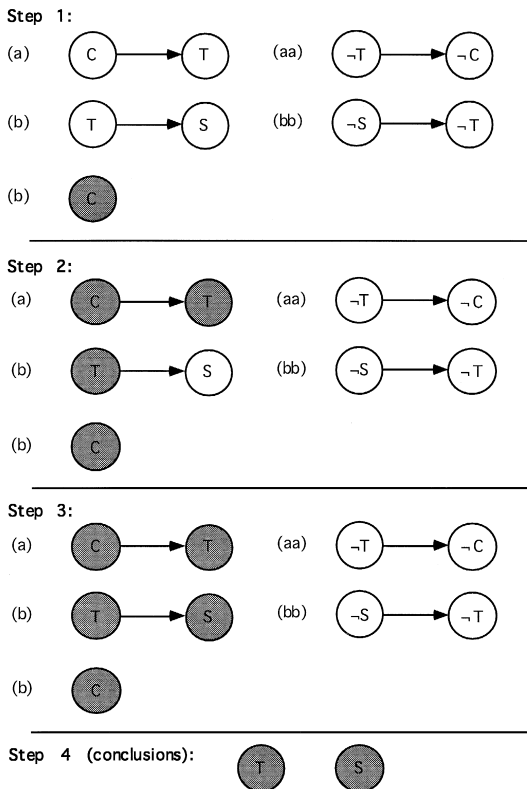


Fig. 10. Example in the training material presented to the experimental group.

is not a circle', and it can also be easily represented in logic graph. (aa) in step 1 depicts the logic graph representation of this statement. The node ' $\neg T$ ' represents 'there is not a triangle', and the node ' $\neg C$ ' represents 'there is not a circle'. The symbol ' \neg ' represents the meaning of 'not'. Similarly, (b) in step 1 depicts the representation of statement 2—'If there is a triangle ('T') then there is a square ('S')', and (bb) depicts its equivalent statement—'if there is not a square (' $\neg S$ ') then there is not a triangle (' $\neg T$ ')'. (c) in step 1 depicts the representation of statement 3—'there is a circle'. We shade node 'C' to indicate that 'C' is a given fact.

Step 2 illustrates how to make inferences with the logic graph method. We first shade all nodes labeled with 'C' in the graph. Then we shade all the nodes reachable from node 'C' (by traversing the arrows). Thus, node 'T' is shaded. This operation (of traversing the arrows) represents the action of making an inference based on the statements 'if 'C' then 'T''

and 'C'. We will continue the process of shading all reachable nodes by traversing the arrows until we can't go any further. In this example, we can reach node 'S' from node 'T' by traversing arrows. Therefore, 'S' is shaded (Step 3). Since we can't go any further from node 'S', we stop at this step. The conclusions that we can draw are represented by all of the shaded nodes (Step 4).

Then, we can check the list of the choices and circle the correct conclusions. Since we have node 'T' shaded in example 1, we choose the conclusion (c) 'there is a triangle' (Node 'S', which represents 'there is a square', is also shaded; but, it is not in the list).

5.4. Test problems

The test problem booklet consisted of 15 problems. Each problem contained three conditional statements and one simple statement. In addition, the subjects in the experimental group were provided with a logic graph representation of the statements. The task for the subjects was to evaluate the given statements and choose all of the correct conclusions that deductively followed from the given statements.

The conditional statements in the problems are given in either 'if-then' form or 'only-if' form. The following is a standard 'if-then' statement.

If John is a lawyer, then Bill is not a doctor.

In this statement, 'John is a lawyer' is the precondition, and 'Bill is not a doctor' is the conclusion. According to the rules of logic, if the precondition of a statement is satisfied, its conclusion must follow. For instance, if we have 'John is a lawyer' as a fact, we must conclude 'Bill is not a doctor.' This type of deductive reasoning is formally called modus ponens in logic.

Every conditional statement has a contrapositive equivalent statement, which is just a different way to state the original statement. For instance, the above statement, 'if John is a lawyer, then Bill is not a doctor' is equivalent to 'if Bill is a doctor, then John is not a lawyer.' Therefore, if we have 'Bill is a doctor' as a fact, we must conclude 'John is not a lawyer'. This type of deductive reasoning is called modus tollens in logic.

Johnson-Laird and Byrne have shown in their experiments that reasoning with modus tollens is more difficult than the reasoning with modus ponens

Table 1
Test problem design

Problem ID	Problem type
#301	Three modus ponens statements
#302	Three modus ponens statements
#303	Two modus ponens and one modus tollens statements
#304	Two modus ponens and one modus tollens statements
#305	Two modus ponens and one modus tollens statements
#306	One modus ponens and two modus tollens statements
#307	One modus ponens and two modus tollens statements
#308	One modus ponens and two modus tollens statements
#309	Three modus tollens statements
#310	Three modus tollens statements
#311	One modus ponens and two modus tollens (with one ‘only-if’) statements
#312	One modus ponens and two modus tollens (with one ‘only-if’) statements
#313	One modus ponens and two modus tollens (with one ‘only-if’) statements
#314	One modus ponens and two modus tollens (with one ‘only-if’) statements
#315	Three modus tollens (with two ‘only-if’) statements

(see Ref. [25], p. 55). According to the mental model theory, people do not follow the logical rules in their deductive reasoning processes, but rather represent conditional statements as mental models. Because reasoning with modus tollens involves more complicated manipulation of mental models than reasoning with modus ponens, people tend to make more mistakes in the former type of reasoning.

A conditional statement can also be in an ‘only-if’ form. For instance, the statement, ‘John is a lawyer only if Bill is not a doctor’, is equivalent to ‘if Bill is a doctor, then John is not a lawyer’, or ‘if John is a lawyer, then Bill is not a doctor’. Although it is easy to transform an ‘only-if’ statement into an ‘if-then’ statement, Johnson-Laird has shown that people tend to make more mistakes with ‘only-if’ statements than with ‘if-then’ statements when they make inferences. Johnson-Laird also provides an explanation of this in terms of his mental model theory.

Our purpose in this experiment was to investigate the effectiveness of the logic graph method in various types of problems. We designed the test problem with different levels of difficulty by combining the modus ponens inferences and modus tollens inferences, as well as, ‘if-then’ statements and ‘only-if’ statements. Our design of the test problem is slightly different from Johnson-Laird’s approach. In Johnson-Laird’s experiments, the test problems were given as a single conditional statement, then subjects

were asked to derive any conclusions that follow the statement. In our experiment, the test problems contain three conditional statements, and the answers were presented in multiple choice format. In general, our test problems are more difficult than Johnson-Laird’s. Although our experimental results generally agree with what the mental model theory predicted, we found some differences. We believe that these differences are caused by different forms of solutions given in the multiple choice. We will analyze the experimental results below. In our experiment, the order of the test problems in the test booklet was randomly assigned for each subject. Table 1 lists the test problems by type. The complete set of test problems is in Ref. [23].

5.5. Results and discussion

Table 2 presents the percentage of subjects who made at least one error for each problem.¹⁶ As

¹⁶ Some 57 subjects participated in the experiment, and 54 subjects completed the test (27 subjects in each group). One subject in the control group and two subjects in the experimental group didn’t answer all of the questions. The experimental result analysis shown in Table 2 does not include these incomplete questionnaires. However, according to our worst case scenario analysis, removal of these three questionnaires does not significantly affect the analysis results shown in Table 2.

shown in Table 2, the subjects in the experimental group (group B) generally performed better than the subjects in the control group (group A). The subjects in group B made fewer errors than the subjects in group A did on 13 test problems out of the total 15 test problems. Among those 13 problems, the subjects in group B significantly outperformed the subjects in group A on nine test problems ($p < 0.05$). In the other two test problems, the subjects in both groups had the same performance in one case (id #301, the subjects in both groups made the equal number of errors on one problem), and the subjects in group A performed slightly better than the subjects in group B did in the other case (id #311).

The test problems were designed to be at different levels of difficulty. We observed that the subjects in both groups performed quite differently across the test problems. For example, in the control group, only 4% of the subjects made at least one error on problem #301, and 82% of the subjects made at least one error on problem #304. This difference is

Table 2

Percentage of the subjects who made at least one error for each test problem and Z-test comparison between the control group and the experimental group

Problem ID	Percentage of errors in control group (A)	Percentage of errors in experimental group (B)	Z-test
#301	4 ± 7	4 ± 7	Z = 0, $p = 0.5$
#302	19 ± 15	15 ± 13	Z = 0.365, $p = 0.358$
#303	22 ± 16	19 ± 15	Z = 0.338, $p = 0.368$
#304	82 ± 15	30 ± 17	Z = 3.834, $p < 0.001$
#305	19 ± 15	4 ± 7	Z = 1.732, $p < 0.05$
#306	70 ± 17	26 ± 17	Z = 3.628, $p < 0.001$
#307	11 ± 12	0 ± 0	Z = 1.728, $p < 0.05$
#308	7 ± 10	4 ± 7	Z = 0.594, $p = 0.276$
#309	19 ± 15	11 ± 12	Z = 0.766, $p = 0.222$
#310	78 ± 16	30 ± 17	Z = 3.548, $p < 0.001$
#311	11 ± 12	15 ± 13	Z = -0.41, $p = 0.675$
#312	40 ± 19	4 ± 7	Z = 3.273, $p < 0.001$
#313	37 ± 18	7 ± 10	Z = 2.619, $p < 0.001$
#314	70 ± 17	26 ± 17	Z = 3.628, $p < 0.001$
#315	26 ± 17	7 ± 10	Z = 1.826, $p < 0.05$

The numbers of percentage errors in the table are the means of percentage errors in each type of the problems in the corresponding groups. Error margins indicated in the table are half-widths of 95% confidence intervals for the respective conditions.

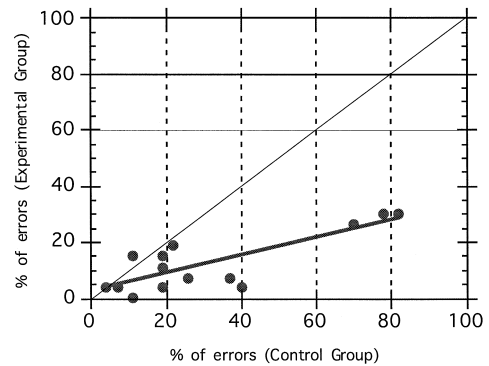


Fig. 11. The performance comparison of the control and experimental groups across the test problems.

quite significant ($Z = 5.78$, $p < 1 \times 10^{-8}$). In general, the difference reflects the level of difficulty in the test problems in terms of the two types of reasoning processes (modus ponens and modus tollens) and two types of conditional statements ('if-then' and 'only-if'). Johnson-Laird has shown that, in 'if-then' conditional statements, people make more mistakes with modus tollens than with modus ponens. He has also shown that, in the 'only-if' conditional statements, there is no significant difference between modus ponens and modus tollens types of reasoning. But people are more likely to have problems with 'only-if' statements than 'if-then' statements [25]. Generally, our experimental results are consistent with those of Johnson-Laird.

Fig. 11 illustrates the performance comparison of the control and experimental groups across the test problems. Each point in the figure corresponds to a test problem with the percentages of errors made by the subjects in the control group (horizontal axis) and the experimental group (vertical axis).¹⁷ Since almost all the corresponding points in Fig. 11 fall below the diagonal line, it concisely shows that the subjects in the experimental group performed better than the subjects in the control group. The trend of improvement in the experimental group is shown by a simple linear regression. On the easy problems (the

¹⁷ There are 15 points in the figure, and two points (problem #306 and problem #314) are overlapped.

percentage of errors is less than 25% in the control group), the effectiveness of the logic graph method is not apparent. As the problem difficulty level increases, the improvement becomes more and more manifest. The overall performance analysis was done with a Kruskal–Wallis test. The test results ($H = 5.395$, $\alpha < 0.05$) demonstrate that the subjects in the experimental group performed significantly better than the subjects in the control group. See Ref. [23] for a detailed discussion of the findings.

In summary, our experimental results indicate that the logic graph representation method is a useful tool for helping people in deductive reasoning. Through the graphical representation of conditional statements, deductive reasoning processes are simplified as graph traversal. This procedure is easily understood and followed. In theory, no matter what forms ('if–then' or 'only–if') of conditional statements are given, no matter what types of reasoning processes (modus ponens or modus tollens) are required, people only need to follow one simple rule to manipulate the graph to reach correct conclusions. However, in our experiment the subjects still performed differently across the test problems. Our explanation is that people have a tendency to rely on their mental models in their reasoning processes. By examining the test booklets of the experimental group, we found that a large number of mistakes were made because the subjects didn't complete 'coloring' the graph before making their choices, especially when the conclusion appeared to be obvious. In these cases, they just chose the 'correct' conclusions based on their mental models. This explains our observation that the pattern of the error rates in the experimental group (B) is similar to that of the control group (A). In fact, when people understand and follow the logic graph method in deductive reasoning, they have little chance to make mistakes. In the experimental group (B), 37% of the subjects made 0 errors, and 56% of the subjects made one error or less. In the control group (A), however, no subjects made one error or less, only 13% of the subjects made two errors, and 87% of subjects made three or more errors. Although our experiment has been limited on testing deductive reasoning based on conditional statements, we believe our conclusion can be generalized—that is, with appropriate supporting tools, people can significantly improve their ability in reasoning.

6. Conclusion

This has been an exploratory exercise. The concept of an argumentation DSS, particularly one with hypermedia elements, is of recent vintage. Neither the requirements nor the operating characteristics of such systems are well understood. Our implementation is very much a prototype. The inferences it makes are elementary and in many ways naïve (e.g., in realistic systems it would not be reasonable to present all known documents that mention presidential polling). Still, the prototype makes inferences and supports hypertext. If, as we have demonstrated, relevant textual documents can be found and displayed using node-specific contextual information, then so can documents in media other than text. There is no reason to think that the principles and the architecture will not generalize, once appropriate indexing schemes and bridge laws [5] are developed for multimedia documents. Finally, our experiment indicates much promise for finding technology that effectively supports reasoning and argumentation of a logical sort. But it is only one experiment and is limited to deductive inference. Very much remains to be investigated with respect to these systems.

A problem that exercises us a great deal is whether very many people will find it useful or attractive to work with an argumentation system, regardless of how friendly and clear it is. We suspect that philosophers and lawyers, who work professionally with arguments, are promising users, but there are serious questions whether many others can be induced to acquire the habits of mind required to reason carefully with arguments. Who, however, would have predicted 30 years ago the widespread, personal, interactive use of computer software? There is reason to hope that, with the accumulation of experimental knowledge, we may yet learn how usefully to bring hypermedia argumentation DSSs to a broad audience.

References


- [1] K.D. Ashley, *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*, The MIT Press, Cambridge, MA, 1990.

- [2] J. Baron, Review of the Skills of Argument by Deanna Kuhn, *Informal Logic* XIV (1) (1992) 59–67.
- [3] E. Berk, J. Devlin (Eds.), *Hypertext/Hypermedia Handbook*, McGraw-Hill, New York, NY, 1991.
- [4] M. Bieber, Automating hypermedia for decision support, forthcoming in *Hypermedia*.
- [5] M.P. Bieber, S.O. Kimbrough, On the concept of generalized hypertext, *MIS Q.* 16 (1) (1992) 77–93.
- [6] M.P. Bieber, S.O. Kimbrough, On the logic of generalized hypertext, forthcoming in *Decision Support Syst.*
- [7] R.L. Causey, EVID: a system for interactive defeasible reasoning, forthcoming in *Decision Support Syst.*, 1993.
- [8] E.K. Clemons, MAC—Philadelphia National Bank's strategic venture in shared ATM networks, *J. Manage. Inform. Syst.* 7 (1) (1990) 5–25.
- [9] J. Conklin, *Hypertext: an introduction and survey*, Computer, September 1987, pp. 17–41.
- [10] J. Conklin, M.L. Begeman, gIBIS: a hypertext tool for exploratory policy discussion, *Proceedings of the Conference on Computer-Supported Cooperative Work*, Sept. 1988, pp. 140–152.
- [11] J. Conklin, M.L. Begeman, gIBIS: a tool for all reasons, *J. Am. Soc. Inf. Sci.* 40 (3) (1989) 200–213.
- [12] F.W. Dauer, *Critical Thinking: An Introduction to Reasoning*, Oxford Univ. Press, New York, NY, 1989.
- [13] J.R. Elliott, S.O. Kimbrough, A graph-based theorem prover for sentence logic model management, University of Pennsylvania, Department of Decision Sciences, working paper 90-01-01, 1990.
- [14] E.R. Emmet, *Learning to Think*, Taplinger Publishing, New York, NY, 1980.
- [15] A. Fisher, *The Logic of Real Arguments*, Cambridge Univ. Press, Cambridge, 1988.
- [16] K.A. Frankovic, Reading Between the Polls, *The New York Times*, June 27, 1992.
- [17] D. Gabbay, C. Hogger (Eds.), *Handbook of Logic for Artificial Intelligence and Logic Programming*, Vol. III, Oxford Univ. Press, Oxford, England, 1993.
- [18] K.M. Gallotti, Approaches to studying formal and everyday reasoning, *Psychol. Bull.* 105 (1989) 331–351.
- [19] M.R. Genesereth, N.J. Nilsson, *Logical Foundations of Artificial Intelligence*, Morgan Kaufmann Publishers, Palo Alto, CA, 1987.
- [20] M.L. Ginsberg (Ed.), *Readings in Nonmonotonic Reasoning*, Morgan Kaufmann Publishers, Los Altos, CA, 1987.
- [21] F.G. Halasz, Reflections on note cards: seven issues for the next generation of hypermedia systems, *CACM* 31 (7) (1988) 836–855.
- [22] S.H. Hashim, *Exploring Hypertext Programming: Writing Knowledge Representation and Problem-Solving Programs*, Windcrest Books, TAB BOOKS, Blue Ridge Summit, PA, 1990.
- [23] H. Hua, Theory and applications of argumentation support systems, PhD thesis, University of Pennsylvania, Philadelphia, 1995.
- [24] P.N. Johnson-Laird, *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*, Harvard Univ. Press, Cambridge, MA, 1983.
- [25] P.N. Johnson-Laird, R.M.J. Byrne, *Deduction*, Lawrence Erlbaum Associates, Publishers, Hillsdale, USA, 1991.
- [26] H. Kahane, *Logic and Contemporary Rhetoric: The Use of Reason in Everyday Life*, 5th edn., Wadsworth Publishing, Belmont, CA, 1988.
- [27] D. Kahneman, P. Slovic, A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge Univ. Press, Cambridge, England, 1982.
- [28] D. Kelley, *The Art of Reasoning*, W.W. Norton & Company, New York, NY, 1988.
- [29] S.O. Kimbrough, A graph representation for management of logic models, *Decision Support Syst.* 2 (1986) 27–37.
- [30] S.O. Kimbrough, The argumentation theory of decision support systems, University of Pennsylvania, Decision Sciences Department working paper 87-12-07, 1987.
- [31] S.O. Kimbrough, Notes on the argumentation theory for decision support systems, *Proceedings of the 1990 International Society on DSS Conference*, Austin, TX, September 1990, pp. 17–39.
- [32] S.O. Kimbrough, An introduction to the method of sweeping presumptions for modeling nonmonotonic reasoning, in: J.F. Nunamaker Jr. (Ed.), *Proceedings of the 24th Annual Hawaii International Conference on System Sciences*, Vol. III, DSS and Knowledge-Based Systems and Collaboration Technology Tracks, IEEE Computer Society Press, Los Alamitos, CA, 1991, pp. 339–348.
- [33] S.O. Kimbrough, F. Adams, Why nonmonotonic logic?, *Decision Support Syst.* 4 (1988) 111–127.
- [34] S.O. Kimbrough, H. Hua, On modeling nonmonotonic reasoning with the method of sweeping presumptions, *Minds and Machines* 1 (4) (1991) 393–416.
- [35] S.O. Kimbrough, S.A. Moore, On obligation, time, and defeasibility in systems for electronic commerce, in: J.F. Nunamaker Jr., R.H. Sprague Jr. (Eds.), *Proceedings of the 26th Annual Hawaii International Conference on System Sciences*, Vol. III, Information Systems: DSS/Knowledge-Based Systems, IEEE Computer Society Press, Los Alamitos, CA, January 1993, pp. 493–502.
- [36] S.O. Kimbrough, C.W. Pritchett, M.P. Bieber, H.K. Bhargava, The coast guard's KSS project, *Interfaces* 20 (6) (1990) 5–16.
- [37] D. Kuhn, *The Skills of Argument*, Cambridge Univ. Press, Cambridge, England, 1991.
- [38] H. Kyburg, R. Loui, G. Carlson (Eds.), *Knowledge Representation and Defeasible Reasoning*, Kluwer Academic Publishers, Boston, MA, 1990.
- [39] J. Lee, Sibyl: a qualitative decision management system, in: P. Winston (Ed.), *Artificial Intelligence at MIT: Expanding Frontiers*, 1, The MIT Press, Cambridge, MA, forthcoming.
- [40] J. Lee, T.W. Malone, Partially shared views: a scheme for communicating among groups that use different type hierarchies, MIT Sloan School working paper, SSM WP #2052-88, March 1988.
- [41] J. Lee, T.W. Malone, How can groups communicate when

- they use different languages? Translating between partially shared type hierarchies, MIT Sloan School working paper, SSM WP #3076-89-MS, Sept. 1989.
- [42] D.G. Lowe, Co-operative structuring of information: the representation of reasoning and debate, *Int. J. Man–Machine Stud.* 23 (1985) 97–111.
- [43] D.G. Lowe, SYNVIEW: the design of a system for cooperative structuring of information, *Computer-Supported Cooperative Work (CSCW '86) Proceedings*, Austin, TX, December 3–5, 1986.
- [44] C.C. Marshall, F.G. Halasz, R.A. Rogers, W.C. Janssen Jr., Aquanet: a hypertext tool to hold your knowledge in place, *Proceedings of Hypertext 1991*, San Antonio, TX, 1991, pp. 261–276.
- [45] D. McDermott, J. Doyle, Nonmonotonic logic I, *Artif. Intell.* 13 (1980) 41–72.
- [46] D. Nute, Defeasible reasoning and decision support systems, *Decision Support Syst.* 4 (1988) 97–110.
- [47] D. Nute, Review of readings in nonmonotonic reasoning, *Philos. Psychol.* 2 (1989) 351–355.
- [48] D. Nute, K. Erk, Defeasible logic graphs: I. Theory, *Decision Support Syst.*, this issue.
- [49] D. Nute, Z. Hunter, C. Henderson, Defeasible logic graphs: II. Implementation, *Decision Support Syst.*, this issue.
- [50] R.C. Pinto, J.A. Blair, *Reasoning: A Practical Guide*, Prentice-Hall, Englewood Cliffs, NJ, 1993.
- [51] J. Pollock, Self-defeating arguments, *Minds and Machines* 1 (4) (1991) 367–392.
- [52] M. Reinfrank, J. de Kleer, M.L. Ginsberg, E. Sandewell (Eds.), *Non-Monotonic Reasoning*, Springer, Berlin, 1989.
- [53] R. Reiter, A logic for default reasoning, *Artif. Intell.* 13 (1980) 81–132.
- [54] R.D. Rieke, M.O. Sillars, *Argumentation and the Decision Making Process*, 2nd edn., Scott, Foresman, Glenview, IL, 1984.
- [55] G. Shafer, J. Pearl (Eds.), *Readings in Uncertain Reasoning*, Morgan Kaufmann Publishers, San Mateo, CA, 1990.
- [56] P. Smolensky, B. Bell, B. Fox, R. King, C. Lewis, Constraint-Based Hypertext for Argumentation, *Hypertext '87*, (November 13–15, 1987), pp. 215–246.
- [57] R.M. Stein, Browsing through Terabytes—Wais opens a new frontier in personal and corporate information services, *Byte* 16 (5) (1991) 157–164.
- [58] P.A. Strassmann, P. Berger, E.B. Swanson, C.H. Kriebel, R.J. Kauffman (Eds.), *Measuring Business Value of Information Technologies*, ICIT Press, International Center for Information Technologies, Washington, DC, 1988.
- [59] S.E. Toulmin, *The Uses of Argument*, Cambridge Univ. Press, London, 1958.
- [60] S. Toulmin, R. Rieke, A. Janik, *An Introduction to Reasoning*, 2nd edn., Macmillan, New York, 1984.
- [61] A. Tversky, D. Kahneman, Judgment under uncertainty: heuristics and biases, *Science* 185 (1974) 1124–1131.
- [62] A. Tversky, D. Kahneman, The framing of decisions and the psychology of choice, *Science* 211 (1981) 453–458.
- [63] D.N. Walton, *Informal Logic: A Handbook for Critical Argumentation*, Cambridge Univ. Press, Cambridge, England, 1989.
- [64] D.N. Walton, *The Place of Emotion in Argument*, The Pennsylvania State Univ. Press, State Park, PA, 1992.
- [65] D.N. Walton, *Plausible Argument in Everyday Conversation*, State University of New York Press, State University Plaza, Albany, NY, 1992.
- [66] A. Weston, *A Rulebook for Arguments*, Hackett Publishing, 1987.



Gary H. Hua is a senior architect at Reed Technology and Information Services. He holds a PhD in management information systems from the Wharton School, University of Pennsylvania. Prior to joining Reed Technology, Dr. Hua was Principal Software Engineer at Unisys and an Assistant Professor in Information Systems at the New Jersey Institute of Technology. Currently, he is working in the areas of internet and intranet technology and electronic commerce.



Steven O. Kimbrough is Professor in the Wharton School, University of Pennsylvania. He holds a PhD in philosophy from the University of Wisconsin at Madison. His teaching and research interests include decision support systems, logic modeling, electronic commerce, and computational theory of rationality.