

Reasoned Use of Expertise in Argumentation

DOUGLAS N. WALTON

*Department of Philosophy
University of Winnipeg
Winnipeg R3B 2E9
Canada*

ABSTRACT: This article evaluates the strengths and weaknesses of arguments based on appeals to expertise. The intersection of two areas is explored: (i) the traditional *argumentum ad verecundiam* (literally, "appeal to modesty;" but characteristically the appeal to the authority of expert judgment) in informal logic, and (ii) the uses of expert systems in artificial intelligence. The article identifies a model of practical reasoning that underlies the logic of expert systems and the model of argument appropriate for the informal logic of the *argumentum ad verecundiam*.

KEY WORDS: *argumentum ad verecundiam*, fallacies, expert systems, plausible reasoning, informal logic, logic of dialogue, practical reasoning, expert opinion.

Recently there have appeared to be some indications of interest in the topic of argumentation¹ among researchers in the field of artificial intelligence, and expert systems particularly. Conversely, some scholars in the field of argumentation have begun to adapt AI techniques and results to the concerns of informal logic.² The intersection of these two areas of study is clearly a fruitful field for future research in argumentation, and in this paper one localized but highly important aspect of it will be discussed.

Traditionally, the *argumentum ad verecundiam* or appeal to authority (typically the authority of expert judgment) has been considered to be an informal fallacy. Nevertheless, occasionally it has been recognized that arguments based on expertise, or appeals to expert judgment, can be reasonable arguments for a conclusion on how to act, even if they are often weak and even, in an important sense, "subjective" kinds of arguments, based on reasonable expectations that may admit of exceptions.³ However, the logic of such arguments has always been elusive, for they do not appear to be deductively valid or inductively strong types of argument. With the advent of expert systems however, it has become apparent that arguments based on expertise may have a logic, or at least a coherent basis as a distinct type of reasoned argument.

The evaluation of the strengths and weaknesses of arguments based on appeals to expertise is the subject of this paper.

I. ARGUMENTUM AD VERECUNDIAM

The traditional *argumentum ad verecundiam* (literally, the argument from modesty) is portrayed by logic textbooks as an informal fallacy, namely the illicit use of an appeal to authority in argument. Among the different kinds of authority cited, the most usual and characteristic of the *ad verecundiam* is the authority of expertise.⁴

Among the kinds of abuses cited by the texts under the heading of the *ad verecundiam* fallacy, three stand out. First, an authority may be cited on a question in one field where the authority is an expert only in some other (not closely related) field. Second, there is the case where the authority cited is not really an expert in the relevant field, but only a figure of prestige or glamor, an opinion trend-setter. Third, there are the cases where the appeal is so vague that the "experts" or fields of expertise are not even named or identified.⁵

These cases rightly indicate that the questionable use of appeals to expert authority to influence opinion in argumentation is commonplace, and worth guarding against. But does this suggest that all appeals to expert authority in argument are fallacious? It should not. For taking into account the intelligent and informed opinion of a legitimate expert on a question that one needs to act upon may be reasonable in argument, if no better or more direct evidence is available when the decision must be made.⁶ This approach seems practical and reasonable, in some cases at any rate, and if so, an argument based on an appeal to expertise need not always be fallacious.

The problem then is the following. If arguments from expertise can sometimes be reasonable and other times be fallacious, how can we distinguish between the reasonable and fallacious instances? Practical studies of interesting cases of appeals to expert authority in argumentation in Woods and Walton (1974), Woods and Walton (1982), and Walton (1985), indicate six requirements to be met for an appeal to expertise to be reasonable. First, the judgment put forward by the expert must actually fall within his field of competence. Second, the cited expert must be a legitimate expert, and not merely a celebrity, or someone not an expert. A third factor is the question of how authoritative an expert is, even if he is a legitimate expert in a field. Questions of specialization within fields of expertise are relevant here. Fourth, if several qualified experts have been consulted, there should be some way of resolving inconsistencies and disagreements that may arise.⁷ Fifth, if objective evidence is also available, this should be taken into account. In particular, an expert should be able to back up his opinion, if queried, by citing evidence in his field. The sixth requirement is that the expert's sayso must be correctly interpreted. This requirement may not be easy to meet, for it involves rendering the expert's exact words in clear language intelligible to a layperson.⁸

guage intelligible to a layperson.⁸

The general form of argument from expertise outlined below incorporates all of these six requirements. In this scheme, E is an individual who is said to be an expert source, A is a proposition, and D is a domain of knowledge, a particular field of competence of expertise.

- (AE)
1. E is a genuine expert in D.
 2. E asserts that A.
 3. A is within D.
 4. A is consistent with what other experts say.
 5. A is consistent with available objective evidence.
-

Therefore, A can be accepted as plausible.

The phrasing of the conclusion of the form of the argument from expertise above clearly indicates that (AE) is meant to be a plausible argument, as opposed to a deductively valid or inductively strong type of argument. Is that acceptable?

Hamblin (1970, p. 218) proposed considering a deductive form of the argument from expertise.

- (DE)
1. Everything E says is true.
 2. E asserts that A.
-

Therefore, A is true.

But there is a problem with (DE). Suppose there are two experts, E¹ and E². But suppose E¹ says that A and E² says that $\neg A$. By (DE), we must conclude that A is true and also that $\neg A$ is true. But this is a contradiction. Of course, this outcome would be blocked if we took the fourth premise of (AE) into account in a revised version of (DE). But that move would run into conflict with premise 1. of (DE). For if there is a real possibility of conflict with other experts, then it could not be literally true that *everything* E says is true, for any E. Clearly the schema (DE) rests on the concept of an omniscient expert, something that flies in the face of any very practical approach to the evaluation of argumentation based on appeals to expertise. Thus characteristically, realistic appeals to expertise are not essentially deductive in their essential structure. Practically speaking, appeals to authority are best viewed as inherently fallible.

Salmon (1963, p. 64) proposed a seemingly more practical model of the argument from expertise.

- (IE)
1. The vast majority of assertions made by E are true.
 2. E asserts that A.
-

Therefore, A is probably true.

This form of argument is supposed by Salmon to be inductive in nature. The truth of the premises make the conclusion probable or inductively

strong. But this inductive model shares essentially the same problem that we had above with the deductive model. The reason lies in the negation rule of the probability calculus: the probability of a proposition $\neg A$ is calculated as 1 minus the probability value of A . Where experts flatly disagree then, the possibility arises through (IE) that A could be probably true and also probably false at the same time. The inductive conception is no better off.

These results leave open a third possibility for an argument structure for (AE), namely the kind of account of plausible reasoning given by Rescher (1976) where both a proposition A and its negation $\neg A$ can be plausible in the same case. And indeed, the whole idea of the appeal to expertise as an essentially subjective but sometimes practically reasonable type of argumentation strongly suggests construal of (AE) as an instance of plausible reasoning.

But what is the model of practical reasoning within which arguments based on expertise can be modelled as reasonable kinds of arguments? We return to this fundamental question below. In order to approach a better understanding of the question, however, let us turn to the concept of an expert system. For as we will see, the nature of an expert reasoning process as based on plausible rules of inference that are neither deductive nor inductive, in all cases, has been well recognized in AI.

2. EXPERT SYSTEMS

An *expert system* is a computer program that simulates the expertise of a human expert in a specific domain of knowledge. Expert systems have several distinctive characteristics. First, they work best in some specific domain of knowledge where expertise can be brought to bear, and do not tend to work well outside such a domain, or if the domain is not narrowly focussed. Second, expert systems are meant to work on problems requiring the experience of expert skill and judgment, not on problems where an algorithm or deductive solution exists. Third, an expert system separates the set of facts in the domain from the set of inference rules or reasoning methods used to derive conclusions from the set of propositions designated as facts. Fourth, an expert system should have the capability to give an explanation of its reasoning in a particular instance if queried.⁹

A wide variety of expert systems currently in use have been very successful in a range of applications. The expert system DENDRAL has been widely used to identify chemical structures. The geology expert system PROSPECTOR has been used to predict the location of mineral deposits. The expert system DOC can diagnose computer crashes better than human experts. The expert system PUFF diagnoses pulmonary disorders at the Pacific Medical Center in San Francisco.¹⁰

Rule-based expert systems can use rules of the if ... then (conditional) sort in sequential form where several *modus ponens* steps, for example, can be put together in a chaining sequence using hypothetical syllogism.




Fig. 1.

The resulting graph of an argument in an expert system reasoning sequence bears many important similarities to methods of diagramming complex arguments in informal logic. This technique has been formalized in Shoemith and Smiley (1980) and Walton and Batten (1984). A key feature of expert system reasoning however is the distinction between *forward and backward chaining*. The above example illustrated forward chaining to conclude D. However, given D, we could reason backwards in a backward chaining sequence, and conclude that D is based on fact A.

The reasoning structure of expert systems is also related to the study of circular argumentation in informal logic. As is shown in the software manual of Intellware (1986, p. 30), where a hypothesis appears in a forward or backward chain which is trying to prove that very hypothesis, the expert system reasoning may "loop" forever. Consider the following type of pattern.




Fig. 2.

The problem is that E has been used as a basis for inferring the conclusion G. But then the sequence of inferences "loops back" from G to E, as well. According to the manual, Intellware (1986, p. 31), designers of knowledge bases for expert systems generally try to avoid using rules that could lead to circles, because circular reasoning could increase performance

costs in executing expert programs. Similarly, in Walton and Batten (1984), it is shown how conditions on reasoning can be set that eliminate circular argumentation as "fallacious." However, it is argued in Walton and Batten that circular reasoning is not necessarily fallacious in all cases, but only subject to criticism in certain contexts of reasoning.

Two farther characteristics of expert systems are worth noting here. One is that an expert system needs to have a method of conflict resolution to apply to a case where too many conclusions can be drawn from a given knowledge based by the rules of inference. This too parallels techniques used in informal logic to deal with inconsistencies when too much information is available. See Woods and Walton (1982, chapter 5) and Rescher (1976). In both expert systems and informal logic, different methods of conflict resolution can be employed in different contexts.

Another important characteristic of an expert system is that it must have a user interface, meaning that a non-expert asks the expert system questions, and the system must be able to respond in a helpful way to such questions. Also, the system must also ask questions of the user, because typically the system does not know all the facts. So the system should ask the user to input those facts which the user may know about.¹¹ This means then that the interaction between the system and the user takes the form of question-reply dialogue.

The system does not try to prove the facts supplied by the user, but only asks the user about their truth. Even the fact that some proposition is "askable" can itself be viewed as a fact and added to the system's knowledge base. Hence, each party supplies input through dialogue that generates the reasoning which leads to the conclusion.

The importance of the dialogue interaction between the user and the system is dramatically illustrated by a tutoring expert system called GUIDON developed at Stanford University. GUIDON is a case-method tutoring expert system designed to improve a student's diagnostic ability in medicine.¹² It was designed to teach about the domains of expertise represented in EMYCIN, a generalized diagnostic system based on an earlier system called MYCIN which is used to diagnose infectious diseases and recommend drug treatments. The inference rules of MYCIN did not correspond closely to human expert reasoning, but NEOMYCIN is more natural. GUIDON uses a NEOMYCIN knowledge base, but also uses an overlay student model which represents the knowledge of the student as overlaying sets of understood and misunderstood concepts.¹³ Thus GUIDON is a dialogue structure that works on interactions between the students repository of accepted propositions and the knowledge base of the tutor system. As well as domain rules that come from the knowledge base of the system, GUIDON has tutorial rules which guide discussions, rank domain rules as more or less relevant to a discourse, and rules which modify the knowledge of the student as the dialogue progresses.¹⁴ The precise form of these rules is given in Clancey (1979).

The important thing about an intelligent tutoring (expert) system is that it is built on a model of reasoning as two-way interactive dialogue that takes into account modification of the student's accepted positions in the dialogue in relation to the presentation of the knowledge in the expert system to the student.

3. INFORMAL LOGIC AND EXPERT SYSTEMS

The inference engine and the user interface of an expert system make up the *expert system shell*, which can be applied to the *knowledge base*, the knowledge that is specific to a domain of application. In informal logic, the counterpart to the knowledge base is the *arguer's position*, the arguer's set of commitments, collected together as a set of propositions.¹⁵ In informal logic, the counterpart of the expert system shell are the rules of procedure for dialogue, including rules that regulate the relationship of the participants and control the kinds of moves that can be made, the kinds of questions that can be asked, and so forth.

In many contexts of argument, the arguer's respective commitment-sets are not called "knowledge bases" because they do not necessarily contain a large enough body of well-organized propositions to correspond to some domain of expertise.¹⁶ In reasonable dialogue, it is not generally required in every case that the arguer's position be tightly restricted to one topic, that they be internally consistent, or that they contain conditionals (rules) that are characteristics of what is called a knowledge base in artificial intelligence.

Nevertheless, the basic function of the knowledge base in AI and the arguer's position in informal logic are essentially similar. They both represent a pool of data or "accepted facts" from which the premises of the reasoner are drawn, and against which his inferences and arguments are tested.

Rules in expert systems can be in the form of plausible conditions of the form: *if situation S then action A*. Such rules may have a number attached to the consequent of the conditional, called the "conclusion" of the rule, which indicates the trustworthiness of the rule as applied to any particular situation. The following rule is an example of one in the MYCIN knowledge base.

IF (1) the gram stain of the organism is gram negative and (2) the morphology of the organism is rod, and (3) the aerobicity of the organism is anaerobic, THEN there is suggestive evidence (0.6) that the genus of the organism is Bacteroides.

The number 0.6 given in the conclusion of the rule allows the interpreter of MYCIN to get an idea of how plausible the rule is in any particular case by giving an indication of the scope for possible exceptions. In this way, the user of MYCIN can weigh the advice given by the expert system

alongside his own knowledge of the particular situation, and arrive at his own conclusion.

The conclusion advanced by MYCIN could admit of several kinds of clear exceptions known to the user. Yet even so, as a plausible conclusion, it could be valuable as a partial basis for the user's diagnosis and action.

The theory that an expert's judgment should be taken as a plausible conclusion, rather than as the output of a deductively valid or inductively strong argument, is based on the assumption that to act on expert advice is to act in a situation where better, more direct evidence is not available within the constraints on reasonable action. One has to act on the best reasons within the known information relative to the given situation. Often, in such cases of reasoned action, the burden of proof tilts in a certain direction, given the agent's reasoned estimate of what is likely and possible to expect in a situation of a type familiar to the agent, or familiar to an expert whom the agent has consulted. Such a data base may be based on subjective opinion rather than empirically verified propositions. Even so, by tilting the burden of proof in one direction by plausible reasoning of a sort that may be intrinsically evidentially weak, an inference based on reasonable expectations, and carefully assessed expert advice, may serve as a kind of reasoned argumentation that may carry justified weight in arriving at a conclusion on how to act.

Plausible reasoning has been linked to the logic of dialogue in Hamblin (1970) and Walton (1985) and (1986). But there appears to be a convergence here because it is now well recognized in AI that the logic of the uses of expert systems must also be based on the concept of reasoned twoperson question-reply dialogue.

4. THE LOGIC OF DIALOGUE

The practical uses of expert systems suggest that a reasonable and appropriate requirement is to try for a system that has a rate of success good enough to be useful, even if it does not work very well in some cases. For example, according to Intellware (1986, p. 57) a system that diagnoses computer crashes correctly in only fifty percent of cases, and says "I don't know," when it cannot predict very well, will still be of great help.

Another factor is that the user can and should use his own judgment in interpreting the conclusions put forward by the expert system in a particular instance.

Hence the idea that an expert system only produces a conclusion that is plausible to a certain degree, rather than certain or even highly probable, is highly acceptable to designers and users of expert systems.

Expert systems apart, it makes sense from a point of view of informal

logic that the logic of appeals to expertise reduces to the logic of dialogue. The following proof is outlined from Walton (1986, chapter 7, section 5). Each participant in argument has his set of commitments, his position to defend. But the expert has an extensive and complex position based on his experience and learned skills in a domain of knowledge or craft. As users of this expertise, we as laypersons do not have direct access to this knowledge base. We can only extract it in usable form by a process of intelligent question-answer interaction (reasonable dialogue). Therefore, the logic of appeals to expertise as a source of reasoned argument reduces to the logic of dialogue.

Many of the traditional informal fallacies have been analyzed as correct or incorrect arguments or moves in argument on the model of argument as reasoned two-person question-answer dialogue by Hamblin (1970), Woods and Walton (1982), and others. At first sight, however, it appears implausible that the analysis of the argument from appeal to expertise should require such a two-person interactive model of argument. However, the development of expert systems in AI has very much brought to the forefront the importance of the user interface in relation to the expert knowledge base.

EXPERT SYSTEM ARGUMENTATION

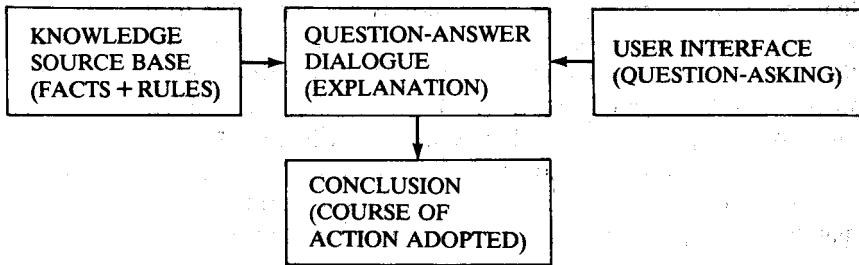


Fig. 3.

Thus contrary to traditional appearances, the use of appeal to expertise in drawing reasoned conclusions does actually require interactive two-person reasoned dialogue as part of the logical structure of the reasoning process.

When asking about the logic of expert reasoning, we need to make the distinction between the internal reasoning of the expert and the use of the expert's conclusion by a second party, usually a layperson in the domain of the request for advice. Even so, the two processes of reasoning must have enough of a common basis for there to be communication between the two parties. The user must ask a question or formulate some goal of the inquiry. The expert must then respond to the question by finding one or more solutions to the goal.

But the process of interaction cannot be that simple in many cases. The expert may have to ask to have the question reformulated before an answer can be given. The user of the information may have to query the expert on what he means in layman's terms.

The internal reasoning carried out by an expert system involves searching in a systematic manner through a knowledge base in order to satisfy the goals set by the question posed. The search procedures used in an expert system language like Prolog may involve some elements of deductive logic, but will have many requirements that are not familiar from deductive logic. Thus to use the system effectively, the user must have knowledge of the language used by the expert system to search out solutions to a goal.

A basic characteristic of programming languages for AI is their goal-oriented structure which involves a hierarchical linkage of actions and goals at different levels of abstraction. Philosophers call this practical reasoning.

5. EXPERTISE AND PRACTICAL REASONING

In searching around for some account of the informal logic of arguments from expertise, Woods and Walton (1974) noted the basic observation that the basis of the expert's conclusion is not accessible to the layperson who has sought his advice. At least, normally it would not be accessible in the form of a set of laws and initial conditions (principles and facts) that would exhibit an implication of the conclusion by deductive (or even inductive) steps of logical inference. Accordingly, the conjecture was raised in Woods and Walton (1974) that expert judgment may be based on intuition, a kind of reasoning derived from familiarity with a subjectmatter of expertise, and susceptible to being spelled out in so many steps of logical reasoning accessible to a non-expert. Even further, Woods and Walton (1974) considered the following possibility: "It is notable that in judgments of expertise the expert's verdict may be based on inarticulable background elements and so not be amenable to total sentential representation."¹⁷ In other words, such judgments could be inaccessible to deductive logic or inductive inference because they may not be even expressible in the form of declarative sentences or propositions. We could call this a strong form of the *inaccessibility thesis* that expert conclusions cannot be tracked back to some set of premises and inference rules (known facts and rules) that yield the basis of the expert judgment.

The inaccessibility thesis tends to be made less plausible in recent times by the advent of machine-based expert systems which have been highly successful in some domains.

One form of the inaccessibility thesis, however, continues to be

advocated. According to Dreyfus and Dreyfus (1986) human expertise is intrinsically different from machine-based expert systems because human experts do not follow rules, except at the novice or beginner level. They cite the example of riding a bicycle (p. 86), a skill learned from experience, but not accessible to us in the form of facts and rules. They cite the case of the boxer (p. 87) who recognizes the moment to attack not by following rules and facts, but by virtue of using practical know-how derived from his training in earlier, similar situations. Dreyfus and Dreyfus (1986) conclude that human expertise is based on intuition, which in turn is based on memories of similar experiences in the past. Thus according to their theory, expertise is based upon a recognition of similarity between new situations and familiar ones - a process of reasoning by analogy. Consequently, they think a clear distinction in principle must be made between "computer rationality" and true expertise.

The Dreyfus case for inaccessibility brings out the important point that analogy is extremely important in expert reasoning, and is equally important in understanding many kinds of reasoning in informal logic. But as indicated by Eliot (1986), current research in AI is tackling the problem of analogical problem-solving as a form of reasoning in expert systems.

Another factor not to be ignored in this connection is the distinction made by Woods and Walton (1972, p. 137) between *special position claims* and *claims from expertise*. A stranger who receives directions from a native of Montreal concerning the location of Jarry Park receives advice from a special position to know that is not a claim based on expertise. Some claims based on expertise may be special position claims, but later on the evidential basis for the claim may become known to the non-expert who solicited the claim. For example, a mathematician's claim to have settled a previously undecided conjecture may be accepted on weaker considerations of authority, but then later settled by checking his proof.

If I have boxed or ridden a bicycle, I may be in a special position to know about these activities *vis a vis* you, who have never tried either of them. I may not be an expert at either of these activities, yet even so my attempts to give you advice on how to approach each of these tasks may be severely limited by your lack of experience. Here, your lack of access to my knowledge may be inherently limited by your own lack of practical experience.

However, the limit to accessibility in such a case is not an inherent limit on expertise because the expertise is based on private intuitions of mine that are not accessible to you. It's just that you lack the requisite practical experiences. All this shows is that some kinds of knowledge are practical knowledge that are based on trial and error. It does not mean that in principle, we could never build a robot to box or ride a bicycle relatively well. It does not show conclusively that expert conclusions are intrinsically inaccessible to a second party. It does indicate that some skills are

essentially practical. But perhaps all that shows is that the informal logic of expert reasoning is a kind of practical reasoning based on a knowledge of actions and practical knowledge in a situation.

Practical reasoning is a kind of goal-directed reasoning of the following general form. A and B are states.

- (P1) I intend to realize A.
 As I see the situation, B is a necessary (sufficient) condition for realizing A.
-

Therefore, I should carry out A.

The above form (P1) is simplified in several respects. First, there may be many steps of action required in a complex (hierarchical) act-sequence. Second, there may be alternative necessary or sufficient steps to consider throughout the sequence. Third, there may be compound goals (with varying priorities). Fourth, there may be side-effects (possible future known consequences of actions) to consider. Fifth, if B is not possible to carry out, in the situation, other means or intentions may be open to consideration.

Practical reasoning has been most notably studied by von Wright (1983). The uses of practical reasoning in informal logic have been studied in Walton (1985). And the theory of rational planning studied by Wilensky (1983) in AI is the natural setting for future studies of practical reasoning.¹⁸

We could sum up our findings as follows. Both the study of the reasoned use of expertise in informal logic and AI require a re-orientation towards the study of practical reasoning as a form of goal-directed plausible reasoning in question-answer two-person dialogue. Practical reasoning is a kind of inference that leads to action in a particular situation and is based on what the agent knows (knowledge base) in a situation as he (she, or it) sees it. Therefore practical reasoning is based on plausible inference and the function of this kind of inference is to shift the burden of proof in reasoned dialogue.

Many questions about the logic of practical reasoning are open to further investigation. At the same time, we have opened up some interesting questions about the logic of expert systems. What does the term 'expert' in 'expert system' really mean? Are human experts intrinsically different from the programs now designated as expert systems? We conclude with a few comments on these open questions.

6. CONCLUDING REMARKS

Human experts may not always arrive at conclusions by the use of rules, and if this is true, human experts may use intuition, or some other form of

information processing that is dissimilar to the machine-reasoning structures used by expert system programs. Whether this is so or not, expert systems are highly interesting to the student of informal logic, because the analysis of *ad verecundiam* argumentation is parallel to studies of the uses of expert systems in many important respects. But how do we define 'expert'? An expert system shell can be applied to any knowledge base. What rule-based program then is not an expert system?

Schank (1982) compares the way an expert stores knowledge about his field to the way a library catalogue does the same job. What is the essential difference? According to Schank, one key difference is that in order to be updated, the library catalogue requires outside intervention, whereas the expert can alter his knowledge as his field or his interests may change. Thus the expert is a self-conscious entity - he knows when he knows something and can change things when he observes something. According to Schank then (1982, p. 1), an expert is a dynamic processor of information that has an awareness of its own internal states and can learn and improve its knowledge internally, by its own internal operations.

But what does "awareness of its own internal states" mean? A program can continually update itself. For example, a financial expert system can modify its conclusions continually as it obtains incoming information on stock quotes. A higher-order hierarchical system can have feedback. It can modify its own internal states, for example, by correcting contradictions in its knowledge base.¹⁹ Does this count as "awareness" or "dynamic information-processing"? Given these questions, it seems best to conclude that there is no clear difference in principle between a human expert and a machine-based expert system. Instead, there are many practical differences of degree that vaguely mark off a constantly changing zone between human expertise and machine expert reasoning.

Thus informal logic has much to learn about the *argumentum ad verecundiam*, and ongoing research in expert systems has a good deal to teach us, by making some of the basic questions more precise in a practical and interesting way.

NOTES

¹ See, for example, Eliot (1986).

² See van Ditmarsch (1986).

³ See Woods and Walton (1982).

⁴ Woods and Walton (1982, p. 97) distinguish between *de facto* appeals to authority (essentially, arguments based on expertise in a domain of knowledge), and *de jure* appeals to authority, which are not based on expertise but on rights invested in institutions in a titular fashion. *De jure* appeals to authority are interesting in their own right, but in this paper the subject will be confined to *de facto* appeals to authority.

⁵ See Hamblin (1970) and Woods and Walton (1982).

⁶ Clements and Ciccone (1984) provide some interesting cases of the use of appeals to expertise in legal argumentation.

- 7 The use of plausibility screening techniques due to Rescher (1976) is carried out in case studies by Woods and Walton (1982, chapter five).
- 8 Techniques for evaluating interpretations of expert judgments are set out in Woods and Walton (1982).
- 9 See Intellware (1986) and *Introductory Readings in Expert Systems*, ed. D. Michie, New York, Gordon and Breach, 1982.
- 10 See P. J. Denning, 'Towards a Science of Expert Systems,' *IEEE Expert*, 1, no. 2, 1986, 80-83.
- 11 See Bratko (1986, p. 315).
- 12 Clancey (1979).
- 13 *Ibid.*, p. 26ff.
- 14 *Ibid.*, p. 28f.
- 15 See Walton (1985).
- 16 See Hamblin (1970) and Walton (1987).
- 17 Woods and Walton (1974, p. 136).
- 18 The analysis of practical reasoning is forthcoming in Douglas N. Walton, *Practical Reasoning*, Totowa, New Jersey, Rowman and Littlefield, 1989. See also D. S. Clarke, *Practical Inferences*, London, Routledge & Kegan Paul, 1985.
- 19 See Wilensky (1983).
- 20 Funding for this research was provided by the following awards: (1) a Killam Research Fellowship from the Killam Foundation of the Canada Council; (2) a Fellowship from the Netherlands Institute for Advanced Study in the Humanities and Social Sciences; and (3) a Research Grant from the Social Sciences and Humanities Research Council of Canada.

REFERENCES

- Bratko, I.: 1986, *Prolog Programming for Artificial Intelligence*, Reading, Mass., AddisonWesley.
- Clancey, W. J.: 1979, 'Tutoring Rules for Guiding a Case Method Dialogue,' *International Journal of Man-Machine Studies* 11, 25-49.
- Clements, C. D. and J. R. Ciccone: 1984, 'Ethics and Expert Witnesses,' *Bulletin of the American Academy of Psychiatry and Law* 12, 127-136.
- Dreyfus, H. and S. Dreyfus: 1986, 'Why Expert Systems Do Not Exhibit Expertise,' *IEEE Expert* 1, no. 2, 86-90.
- Eliot, L. B.: 1986, 'Analogical Problem-Solving and Expert Systems,' *IEEE Expert* 1, no. 2, 17-26.
- Hamblin, C. L.: 1970, *Fallacies*, Methuen, London.
- Intellware: 1986, *Experteach* (software and manual), Intelligence Ware, Inc., Los Angeles.
- Rescher, N.: 1976, *Plausible Reasoning*, Van Gorcum, Assen-Amsterdam.
- Salmon, W.: 1963, *Logic*, Prentice-Hall, Englewood Cliffs.
- Schank, R. G.: 1982, *Dynamic Memory*, Cambridge University Press, Cambridge.
- Shoemith, D. and T. Smiley: 1980, *Multiple-Conclusion Logic*, Cambridge University Press, Cambridge.
- Van Ditmarsch, H. P.: 1986, 'Applications of Abstraction in Argumentation,' *Argumentation: Perspectives and Approaches*, ed. F. van Eemeren, R. Grootendorst, J. A. Blair and C. A. Willard, Dordrecht, Foris Publications, 1987, 162-169.
- Von Wright, G. H.: 1983, *Practical Reason*, Cornell University Press, Ithaca, New York.
- Walton, D.: 1985, *Arguer's Position*, Greenwood Press, Westport, Connecticut and London, England.
- Walton, D.: 1987, *Informal Fallacies*, John Benjamins, Amsterdam.
- Walton, D. and L. Batten: 1984, 'Games, Graphs and Circular Arguments,' *Logique et Analyse* 106, 133-164.

Wilensky, R.: 1983, *Planning and Understanding: A Computational Approach to Human Reasoning*, Addison-Wesley, Reading, Mass.

Woods, J. and D. Walton: 1974, 'Argumentum Ad Verecundiam,' *Philosophy and Rhetoric* 7, 135-153.

Woods, J. and D. Walton: 1982, *Argument: The Logic of the Fallacies*, McGraw-Hill Ryerson, Toronto.