CHAPTER 8

NON-NUMERICAL APPROACHES TO PLAUSIBLE INFERENCE

INTRODUCTION
by
Glenn Shafer and Judea Pearl

Though non-numerical plausible reasoning was studied extensively long before artificial intelligence was developed, the articles in this chapter are all drawn from the artificial intelligence literature.

Our first article, by Allan Collins, discusses how a variety of methods of plausible reasoning might be integrated into a system for responding to queries. Collins' ideas are derived from an analysis of how people respond to simple questions in a teaching environment. They have been developed further in later publications, especially in an article by Collins and Ryszard Michalski (1989).

Our second and third articles, by Raymond Reiter and Terry Winograd, are concerned with the possibilities for adapting deduction to deal with plausible reasoning. Reiter surveys the field of nonmonotonic logic, which has grown up in the 1980s. Nonmonotonic logics hew closely to the standard conception of deductive logic, but they attempt to couple deduction with facilities for retracting plausible assumptions when contradictions arise or new information is obtained. Winograd surveys earlier work, largely from the 1970s, which dealt with the same problem but with less formality and less emphasis on maintaining the primacy of deduction.

Our fourth article, by David Touretzky, is a brief discussion of defaults in inheritance systems. An inheritance system is a graphical scheme for organizing knowledge about classes, their hierarchies, and typical properties of the objects in the classes. Inheritance systems have not been associated closely with uncertain reasoning, but Touretzky's work on incorporating specificity-based arguments into inheritance systems have brought them closer to probability ideas.

Our last article, by Michael Sullivan and Paul Cohen, presents an application of Cohen's theory of endorsements, which was originally developed as a alternative to the probabilistic handling of uncertainty. This article again indicates a convergence between symbolic and numerical approaches; Sullivan and Cohen use numbers to indicate strength of evidence within the endorsement framework.

The remainder of this introduction begins by placing plausible reasoning, and its relation to logic and probability, in a historical perspective. It then looks in more depth on the articles in the chapter, and on the encouragement they give to a rapprochement between AI and probability. Most of the articles in this chapter support, in one way or another, the need in plausible reasoning for structures similar to the conditional-independence structures of probability theory.

A Historical Perspective on Plausible Reasoning

Plausible reasoning—reasoning that leads to uncertain conclusions because its methods are fallible or its premises are uncertain—has a long history in Western thought. From the time of Aristotle, philosophers and logicians have classified and analyzed various types of plausible reasoning, including arguments by analogy, induction, and abduction.

The relation between logic and plausible reasoning has been seen in many different ways. In ancient times, plausible reasoning was sometimes classified under the head of rhetoric, or the art of persuasion, but other writers saw it as part of logic. Many ancient and medieval ancient authors saw Aristotle's syllogism as more basic to logic than plausible reasoning, but the syllogism fell into low repute for many centuries, from around 1500 until the middle of the nineteenth century, and during this period most philosophers put methods of plausible reasoning at the center of their conception of logic (Kneale and Kneale 1962).

Since the late nineteenth century, when modern symbolic logic was invented as a formalization of deduction, methods of plausible reasoning usually have not been allowed to share the name

"logic." Before the advent of artificial intelligence in the late 1950s, however, few proponents of symbolic logic considered it sufficient as a tool for all human reasoning. Symbolic logic was seen as a formalization of deductive reasoning, and deductive reasoning was seen as just one part of reasoning in general. Even in mathematics, where deduction holds formal pre-eminence, it must be supplemented by plausible reasoning (Hadamard 1945, Polya 1954).

In spite of the autonomy of plausible reasoning from logic in modern times, and in spite of wide and continuing interest in plausible reasoning, the topic has never acquired a standard terminology and a core of agreed-on theory in the way that deductive logic has. One reason for this is that formal treatments of plausible reasoning have always tended to be absorbed by probability theory. Philosophers such as Bertrand Russell (1948) insisted that mathematical probability applied to only some instances of plausible reasoning, but their accounts of non-probabilistic plausible reasoning tended to be verbal and impressionistic, not formal. Even George Polya turned to probability theory when he attempted a formal rather than an impressionistic account of plausible reasoning in mathematics.

Mathematical probability was originally intended as a theory of plausible reasoning. The idea of the degree of probability of an opinion was well-established in law and philosophy before mathematical probability was invented, and the attachment of the word probability to the mathematical theory, by scholars such as James Bernoulli in the late 1600s and early 1700s, was an indication of their intention to use the mathematical theory as a general tool for the evaluation of all opinion and all plausible reasoning. Bernoulli entitled his book on probability *Ars Conjectandi*, or the Art of Conjecture.

Throughout the eighteenth century and well into the nineteenth century, the mathematical theory of probability was almost universally seen as the theory of rational belief (Daston 1988). It was only in the mid-nineteenth century that this view began to be discredited by empiricist philosophy, and the frequentist interpretation of probability emerged (Porter 1986). Even after the frequentist interpretation gained an upper hand, philosophers have repeatedly reached to probability theory to unify the ideas of plausible reasoning. For example, Charles Sanders Peirce, the American logician who introduced the term "abduction" into logic, reconciled the frequentist interpretation of probability with the view that probability plays a role in all non-deductive inference. As he put it, "probability is the proportion of arguments carrying truth with them from among any genus" (Feibleman 1970, pp. 123-124).

In this context, the sustained effort within artificial intelligence to develop a formal non-probabilistic and even non-numerical theory of plausible reasoning is quite unique. This effort was made possible by the unusually large role assigned to deductive logic by some of the founders of AI, especially John McCarthy, and by the obvious difficulties involved in computer implementation of probability. Because of these difficulties, most of those who opposed McCarthy and the other logicists saw relatively ad hoc programming, not probability, as the main alternative (Israel 1983).

Plausible Reasoning in AI

The articles in this chapter do not represent all the strands of work on non-numerical plausible reasoning in AI. Aside from the wide-ranging article by Collins, none of these articles deal with analogy or induction. (For references to AI work on analogy, see Winston 1980 and Carbondell 1981; for work on induction, see Holland et al. 1989.) These articles do, however, represent both logicist and non-logicist positions. Reiter is a logicist, while Collins, Winograd, and Sullivan and Cohen can be characterized as non-logicists.

The AI systems described by Winograd tend to use logic-like notation, but as Winograd emphasizes, they do not limit themselves to modus ponens, the primary inference rule of deductive logic. Instead they use many other modes of inference, and they use levels of organization on top of the logical syntax to direct these modes of inference.

Reiter's article provides an excellent survey of nonmonotonic logic, an umbrella term for logicist work on plausible or commonsense reasoning in the 1980s. This work recognizes the need for extended modes of inference, but it attempts to retain a primary role for deduction, and it attempts to retain a semantics based on model theory (see section 7 of Winograd's article).

Sullivan and Cohen's article applies Cohen's theory of endorsements (Cohen 1985) to plan recognition. Endorsements are symbolic representations of different items of evidence, the questions on which they bear, and the relations between them. Endorsements can operate on each other and hence lead to the retraction of conclusions previously reached, but since there is no formal accounting

2

of final conclusions, the process is seen as a procedural implementation of non-monotonic patterns of reasoning rather than as a logic.

Collins's article, though brief, is the most ambitious in this chapter, for it covers the whole range of plausible reasoning, from reasoning by analogy to self-referential meta-reasoning ("I must not have an older brother, because I would know if I did," etc.). As Collins sees it, a plausible reasoning system must integrate all these modes of reasoning and allow their conclusions to reinforce each other or be weighed against each other. In relation to Collins' article, most of the other articles in this chapter are fairly narrow. Nonmonotonic logic, for example, is concerned with only one of Collins's types of inference, meta-inference, or inference based on one's knowledge about one's own knowledge. Winograd reports on a range of programs, but since he is responding to the nonmonotonic formalists, he too emphasizes meta-inference.

Cohen's theory of endorsements is as broad in intention as Collins's theory, but Collins's work is based on a much broader range of examples, and hence his theory deals more thoroughly with the differences among types of inference.

Rapprochement with Probability

The readings in this chapter suggest that current work on plausible and commonsense reasoning in AI can lead to a rapprochement with probability ideas.

At a superficial level, this is most evident in the willingness of the non-logicists to embed numerical probabilities in their systems. Sullivan and Cohen mention in their article that they use numerical weights to represent the strengths of endorsements. Collins emphasizes the need to attach degrees of certainty to inferences, and Collins and Michalski (1989) suggest that these degrees of certainty should be numerical probabilities.

Of deeper significance is the growing recognition, from the viewpoint of almost every formalism, of the need for structures similar to the conditional-independence structures of probability theory (see the article by Pearl, Geiger, and Verma in Chapter 1 and the articles by Pearl in Chapter 6). Reiter, in section 2.3 of his article, emphasizes the need to use model structure, rather than mere rules of thumb, as a basis for diagnosis. Touretzky shows how the proper management of defaults in inheritance networks depends on attention to their topological structure, which is often a conditional independence structure (see Neufeld's article in Chapter 9). Cohen emphasizes that endorsements must be structured so that some operate on others, and this, too, corresponds to the specification of a conditional independence structure.

Perhaps the deepest point of contact between probability and non-monotonic logic revolves around the issue of specificity-based arguments. Both logic and probability must find ways to ensure that inferences be based on the most specific classes for which information is reliable (e.g., the inference that a penguin cannot fly must override the inference that a bird can fly). In the case of default logic, this requires semi-normal rules, which explicitly specify exceptions (e.g., birds fly, unless they are penguins or ostriches or . . .). In the case of circumscription, we must supply priorities among abnormalities (McCarthy 1980). Other approaches, including probabilistic approaches, rely on structure to manage such priorities.

Touretzky argues that the enumeration of exceptions places an impractical burden on the management of inheritance networks, and he shows how attention to "inferential distance" in the network can assure priority for more specific arguments without such explicit enumeration. In the next chapter, we will see that specificity-based priority can also be based on probability theory, even if numerical probabilities are not used, provided that we interpret defaults as statements of high conditional probability, infinitesimally close to one.

Truth maintenance systems, which can be regarded as the practical side of non-monotonic logic, generally give a considerable role to structure, and hence lend themselves relatively easily to integration with probability ideas. One indication of this is the article by Laskey and Lehner in the next chapter, which integrates truth maintenance with Dempster-Shafer theory. Another indication is the article by Pearl in the next chapter, which shows how distinguishing causal from evidential justification, an idea borrowed from probability theory, can enable truth maintenance systems to reason with causation.

The remaining resistance to giving probability a role in nonmonotonic logic may stem from a lack of appreciation of the role of structure in probability theory, and from an exclusive reliance on the frequency interpretation of probability. In section 7 of Reiter's article, for example, we find

probabilistic treatments of plausible inference equated with "statistical readings." We hope that the context provided by this volume of readings will help remedy both these misconceptions.

References.

Bernoulli, James (1713). *Ars Conjectandi*. Basel.

Carbondell, Jaime G. (1981). A computational model of problem solving by anaology. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, Vancouver, pp. 147-152.

Cohen, Paul R. (1985). *Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach*. Pitman: Boston.

Collins, Allan, and Ryszard Michalski (1989). The logic of plausible reasoning: A core theory. *Cognitive Science* **13** 1-49.

Daston, Lorraine (1988). *Classical Probability in the Enlightenment*. Princeton University Press.

Feibleman, James K. (1970). *An Introduction to the Philosophy of Charles S. Peirce*. MIT Press.

Hadamard, Jacques (1945). *The Psychology of Invention in the Mathematical Field*. Dover.

Holland, John H., et al. (1989). *Induction: Processes of Inference, Learning, and Discovery*. MIT Press.

Israel, David (1983). The role of logic in knowledge representation. *IEEE Computer* October 1983, pp. 37-41.

Kneale, William C., and Martha Kneale (1962). *The Development of Logic*. Oxford.

McCarthy, John (1980). Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence* **28** 89-116.

Polya, George (1954). *Mathematics and Plausible Reasoning (Volume I. Induction and Analogy in Mathematics; Volume II. Patterns of Plausible Reasoning)*. Princeton University Press.

Porter, Theodore M. (1986). *The Rise of Statistical Thinking: 1820-1900*. Princeton University Press.

Russell, Bertrand (1948). *Human Knowledge: Its Scope and Limits*. Simon Schuster: New York.

Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the Asociation of Computing Machinery* **23** 689-703.

Articles for Chapter 8

1. Collins, Allan (1978) Fragments of a theory of human plausible reasoning. In D. Waltz, ed. *Theoretical Issues in Natural Language Processing II* (pp. 194-201). University of Illinois.
2. Reiter, Raymond (1987). Nonmonotonic Reasoning. *Annual Review of Computer Science*, Vol. 2, pp. 147-186.
3. Winograd, Terry (1980). Extended Inference Modes in Reasoning by Computer Systems *Artificial Intelligence* 13, pp. 5-26.
4. Touretsky, David S. (1984) Implicit ordering of defaults in inheritance systems. *AAAI-5*. 322-325. (Reprinted previously in Ginsberg's *Readings in Nonmonotonic Reasoning.*)
5. Sullivan, Michael, and Paul R. Cohen (1985). An Endorsement-Based Plan Recognition Program. *IJCAI-85*, 475-479.