

ON USING HUMAN NONMONOTONIC REASONING TO INFORM ARTIFICIAL SYSTEMS

Marilyn FORD

Griffith University, Nathan, Queensland, Australia

People seem adept at drawing tentative conclusions when premises do not lead to a necessary conclusion. In contrast, the artificial nonmonotonic reasoning systems that have been developed are complex and do not function with ease. This apparent difference between human and artificial computational reasoning is sometimes considered puzzling and frustrating - if people can do it so easily, why can't we get computers to do it easily? The present paper explores the ways in which people attempt to solve nonmonotonic problems which contain conflict and shows that people do not in fact reason about these nonmonotonic problems so easily; they jump to conclusions easily, but they do not reason so well. However, some people do manage to sometimes reason quite well, in that their reasoning is based on ideas that are (classically) logically justifiable. This paper explores differences between these reasoners and others who cope less well. It also explores how the identification of the way in which these people reason can be used to inform artificial nonmonotonic reasoning systems.

It is well known in the literature on human reasoning that people are not adept at deductive reasoning (Evans, Newstead, & Byrne, 1993; Johnson-Laird & Byrne, 1991; Ford, 1995). It is often asserted that this stands in stark contrast to our everyday commonsense reasoning (Oaksford & Chater, 1998; Pelletier & Elio, in press). In everyday reasoning, much of the information we deal with is uncertain and does not cover all cases and so the required reasoning often cannot be deductive. The reasoning is often nonmonotonic in that, unlike deductive reasoning, the number of conclusions drawn will not necessarily increase as the number of premises increases - as we get more information, some conclusions might need to be retracted (for overviews, see Ginsberg, 1987; Lukaszewicz, 1991; Antoniou, 1997). Throughout the history of nonmonotonic reasoning in Artificial Intelligence (AI), it has been assumed that people, unlike artificial systems that have been developed, cope well with such reasoning (Brachman, 1990; Kraus, Lehmann, & Magidor, 1990; Gabbay, 1994). Psychologists seem to agree with the idea that people cope well:

Information and Communication Technology, Applied Cognitive Neuroscience Research Centre, Griffith University, Nathan, Queensland, Australia.

Marilyn Ford, Information and Communication Technology, Applied Cognitive Neuroscience Research Centre, Griffith University, Nathan, Queensland, Australia, 4111. E-mail: m.ford@griffith.edu.au

In deductive reasoning, to be sure, human performance is extraordinarily poor and brittle, and only very minute problems can be tackled ... Yet this stands in direct contrast to the case of commonsense reasoning, where we appear to be able effortlessly to recruit vast amounts of knowledge in drawing plausible conclusions (indeed, the entire knowledge base may be in play, rather than two or three premises).

(Oaksford & Chater, 1998, p. 134)

It is because ordinary people so cleverly and effortlessly use default reasoning to solve interesting cognitive tasks that nonmonotonic formalisms were introduced into AI ...

(Pelletier & Elio, in press, Abstract)

Pelletier and Elio (1997, in press) go further and argue that nonmonotonic reasoning is *psychologistic*: that is, they argue that unlike deductive reasoning, where there is an “external standard of correctness”, for nonmonotonic reasoning the “correct” answers should be defined by what “ordinary people” (as a whole) do. They urge the AI community to consider the reasoning of “ordinary people”, saying that data from empirical investigations of “ordinary people” is the very data that formal theories of nonmonotonic reasoning should cover.

The present paper considers the question of how, and to what extent, research on human nonmonotonic reasoning can help inform artificial systems. In doing so, the paper discusses some of the classic problems considered by the nonmonotonic reasoning community in AI and the answers considered as desirable by this community. It then considers some data on human nonmonotonic reasoning.

The AI community on nonmonotonic reasoning

As Pelletier and Elio (1997; in press; Elio & Pelletier, 1993, 1994) note, AI researchers do not gather data on human reasoning when attempting to formalize their theories of nonmonotonic reasoning, but instead use their own intuitions. It seems it is assumed that for many problems the “correct” answer is obvious. A typical problem considered when showing the obviousness of a “correct” conclusion in nonmonotonic reasoning is “Tweety the Penguin”, given in (1) (Touretzky, 1984; Horty, 1994), where the “correct” answer is that Tweety does not fly.


(1) All penguins are birds.

Birds usually fly.

Penguins usually do not fly.

Tweety is a penguin.

Does Tweety fly?



The diagram is presented here simply to help the reader see the structure of the problem. In this and in other diagrams, an arrow that is not dashed represents a strict rule: all of the X are Y. A dashed arrow represents non-strict rules, something like: X are usually/typically/normally/mostly Y. A dashed line that is crossed indicates a non-strict rule with a negative meaning, something like: X are usually/typically/normally/mostly not Y. It should be pointed out that for this classic Tweety example with bird content and for other problems in the literature, it is the form of the arguments that is considered important and not the content used. In this classic case, the structure given is such that it cannot be logically concluded that Tweety either does or does not fly. There is no deductively correct answer. However, in the literature on non-monotonic logic, the preferred conclusion is that Tweety does not fly. The preferred conclusion for this and many other examples is based on what is known as the specificity principle. Although precise interpretations of the principle differ, they have their origin in Touretzky's (1984) notion of inferential distance ordering, which says basically that when there is conflict, information stemming from a subclass should override information stemming from its superclass, with A being a subclass of B "iff there is an inheritance path from A to B" (e.g., Bacchus, 1989; Horty, 1994; Nute, 1994; Schlechta, 1997; and Stein, 1989). In (1), penguin is a subclass of bird, and so information stemming from penguins overrides information stemming from birds. However, this classic problem is not really a good example for showing how easy nonmonotonic reasoning is for people. We all know that penguins never fly in any sense related to wing flapping (and they probably do not fly in planes too often either!) and so we can answer the problem merely by retrieving this fact from memory. In fact, the structure is misleading: the negative non-strict arrow should be strict.

The "Pennsylvanian Dutch" problem, given in (2), is another classic (Horty & Thomason, 1988).

- (2) All native speakers of Pennsylvanian Dutch are native speakers of German.


Native speakers of German are usually not born in the US.

Native speakers of Pennsylvanian Dutch are usually born in Pennsylvania.

All people born in Pennsylvania are born in the US.

Hermann is a native speaker of Pennsylvanian Dutch.

Was Hermann born in the US?



By the specificity principle, the preferred answer is that Hermann was born in the US: there is a conflict stemming from G (native speaker of German) and PD (native speaker of Pennsylvanian Dutch), but PD is a subclass of G, so let the information from PD override the information from G. Many people would give the expected answer. However, many people probably already know that speakers of Pennsylvanian Dutch are normally born in the US and they can quite easily ignore the other side of the argument.

Pelletier and Elio's (1997; in press; Elio & Pelletier, 1993, 1994) call for more studies on human nonmonotonic reasoning and their claim that nonmonotonic reasoning is psychologicistic should be seen as separate issues. They will thus be treated separately in the second and third sections. The question of how people actually cope with nonmonotonic reasoning problems is treated in the fourth section. The issue of whether any positive insights can be gained from human reasoning that could be used in artificial systems is considered in the fifth section.

Reasons for agreeing with Pelletier and Elio on the need to study human nonmonotonic reasoning

There are five main reasons why one might agree with Pelletier and Elio (1997; in press; Elio & Pelletier, 1993, 1994) on the need to study human nonmonotonic reasoning:

- a. The classic problems, such as (1) and (2), are often biased to already known answers and sometimes contain non-strict statements when the content is actually strict, thus eroding one's confidence that the intuitions of AI researchers about preferred responses would hold for the argument forms *per se*. (For other classic problems, see Touretzky, Horty, & Thomason, 1987; Stein 1989).
- b. Where the intuitions of AI researchers conflict, the intuitions of "ordinary people" might show a preference that could be used to settle the dispute.
- c. "Ordinary people" might reason in a way that is different from that con-

sidered by AI researchers and thus give new insight into nonmonotonic reasoning.

- d. Some results obtained by Pelletier and Elio give an insight into some factors that influence human nonmonotonic reasoning and that lead to plausible conclusions; for example, if it is known that some object is an exception to a rule and there is another object that is similar to it, people are more likely to conclude it is also an exception.
- e. In everyday reasoning, we are faced with situations where answers cannot be validly deduced, but where we need to draw tentative conclusions that we can later retract when given more information; nonmonotonic reasoning is thus an important topic, a fact now being recognised by more and more psychologists (see, for example, Schurz, 2002; Pfeifer & Kleiter, 2003, in press, 2005; Dieussaert, Ford, & Horsten, 2004; Benferhat, Bonnefon, & Da Silva Neves, in press; Ford, in press).

Reasons for being sceptical of psychologism

While there are good reasons for studying human nonmonotonic reasoning, there are reasons to be sceptical of psychologism:

- a. People are known to be bad at deductive reasoning, a fact that should make one suspicious of any claim that assumes they are good at nonmonotonic reasoning.
- b. Although there are no conclusive answers for nonmonotonic reasoning problems, some answers may be based on flawed assumptions or erroneous notions. Take (1), for example. Someone might argue that the positive path is stronger because it contains a strict rule, but this is fallacious, a point that is discussed further below. If people give responses based on such flawed assumptions or erroneous notions, then their conclusions should be treated with suspicion.
- c. While people may be quite adept at drawing conclusions quickly when given uncertain and insufficient information and good at retracting their conclusions if necessary, this does not show that they are good at reasoning. It may simply mean they are happy to jump to tentative conclusions until they come across some disconfirming evidence. However, the aim of AI systems is to reach the best conclusion immediately, if there is one; the best conclusion being one that is favoured on consistent, coherent, principled grounds.

So how do people fare with nonmonotonic reasoning problems?


For simple problems, people are certainly happy to draw conclusions that are in keeping with what many AI researchers would expect. Consider, for example, (3) and (4).

- (3) All members of the plant species zillo are small.

Small plant species are usually found in deserts.

Garffi is a zillo.

Is Garffi found in deserts?



- (4) Trendors are usually green animals.

Green animals are usually found on cliffs.

Stordy is a trendor.

Is Stordy found on cliffs?



There is no deductively correct answer for these problems. However, in the nonmonotonic reasoning literature it is assumed that people would want to tentatively conclude in the affirmative, until they receive information to the contrary. Thus, for example, given *all echidnas are mammals, mammals usually do not lay eggs, Susie is an echidna*, we can assume (erroneously in fact) that echidnas do not lay eggs and that Susie does not lay eggs, until we hear otherwise. Ford and Billington (2000) gave 19 university students a series of problems like (3) and (4) (though without the diagrammatic representations) where the content was about fictitious plants and animals and where conclusions could therefore not be simply retrieved from memory. For (3), all 19 subjects concluded that Garffi was likely to be found in deserts; for (4), 17 concluded Stordy was likely to be found on cliffs. These subjects are thus willing on these simple problems to draw conclusions that the literature on nonmonotonic logic would expect.


For problems involving conflicting arguments, people fare less well (Hewson & Vogel, 1994; Vogel, 1996; Ford & Billington, 2000; Ford, in press). Hewson and Vogel (1994; Vogel, 1996) found that most people were reluctant to come to any conclusion for the problems they gave. Ford and Billington (2000) showed that when conclusions could not be retrieved from background knowledge, subjects did not overwhelmingly give the expected answers based on the specificity principle. Thus, for example, in two different studies, only 7 out of 19 and 2 out of 12 subjects drawn from a student population gave the expected, negative, answer to problems with the structure of (1). In one study, only 6 out of 19 gave the expected, positive, answer to problems with the structure of (2). This should not be taken as indicating that the answers expected due to the specificity principle in the nonmonotonic reasoning literature are unjustified. Ford and Billington found five negative factors (N1 – N5) in people's reasoning with nonmonotonic problems and Ford (in press) who, unlike Ford and Billington, used problems that concerned primarily real world items, found a further three (N6 – N8):

N1. Unwillingness to draw a tentative conclusion when faced with conflict and non-strict rules. Perhaps N1 is reasonable given that there are no valid

answers to the problems, but avoiding tentative conclusions would defeat the purpose of trying to give the most reasonable answer in artificial systems.

N2. Counting up the perceived number of relevant positive and negative paths, even when some paths did not actually exist. N2 suggests a misunderstanding of an argument. Consider (5), which gives the structure of a problem studied by Ford and Billington (2000).

(5)




Claiming, as a number of subjects did, that the individual i who is a B was probably a D because there are two positive paths from B to D is simply wrong; there is only one positive path from B to D.

N3. Using path length to respond, regardless of the ordering of rule types. N3 leads to conflicting results depending on whether shorter or longer paths are preferred. Moreover, it shows a lack of understanding that $A \dashrightarrow B \rightarrow C$ is logically equivalent to $A \dashrightarrow C$. That is, from $A \dashrightarrow B \rightarrow C$ it can be deduced that $A \dashrightarrow C$ and path length is thus irrelevant.

N4. Giving preference to an argument that contains the universal quantifier "all". N4 also shows a lack of understanding; the presence of the universal premise in $A \rightarrow B \dashrightarrow C$ or $A \dashrightarrow B \rightarrow C$ simply does not make the argument stronger than $A \dashrightarrow C$. Imagine a database that includes information equivalent to (6).

(6)



One would hope that rather than giving preference to the positive path $A \dashrightarrow B \rightarrow C$, the system would recognize that there must be a problem with

the database as it contains *contradictory* information because $A \dashrightarrow B \rightarrow C$ is logically equivalent to $A \dashrightarrow C$.

N5. Giving preference to either “usually not” over “usually” or “usually” over “usually not”. N5 has no basis and, like N3, leads to conflicting answers, depending on whether “usually” or “usually not” is given preference.

N6. Giving preference on the basis of opinions about the world and not on the basis of the argument. N6 disregards argument structure, which a system probably should not do. If the arguments are not considered, then there is no point putting them in the system.

N7. Giving preference to information stemming from the larger of two groups. N7 confuses size of group with the probability that an individual in the group has a property typical of the group.

N8. Taking a non-strict argument as evidence for the other side of the argument. “X are usually Y” is seen as evidence for the existence of some X that are not Y and is thus seen as supporting the opposing side of the argument. Alternatively, “X are usually not Y” is seen as evidence for the existence of some X that are Y and thus seen as supporting its opposing side. This is slightly different from N5, because for N8 a negative argument is taken to support a positive argument and vice versa. N8 may happen when people do not see that the same type of “reasoning” could apply to both sides of the argument.

A ninth factor should also be mentioned. When interviewing subjects, Ford (Ford & Billington, 2000; Ford, in press) found that they sometimes seemed to be ignoring some of the given information. In such cases, she asked subjects what they were doing with the information they seemed to be ignoring. This is because she wanted the subjects to reason taking all of the information into account. In most of these cases, subjects said they had forgotten about a line and then proceeded to reason trying to take all the information into account. This leads to N9:

N9. Forgetting to take some information into account. Again, there is no sense putting such a thing in an artificial system.

The negative factors N1 - N9 suggest that for problems involving conflicting arguments, one needs to be very wary of relying on people’s reasoning about nonmonotonic problems; it thus seems that nonmonotonic reasoning should *not* be considered psychologistic and that the nonmonotonic reasoning of “ordinary people” (as a whole) cannot be used to help inform artificial systems.

Can any positive insight be gained from people's nonmonotonic reasoning about relative complex problems?

Although Ford and Billington (2000) and Ford (in press) found that people have difficulty with nonmonotonic reasoning, they identified three positive factors that sometimes influence reasoning:

P1. Recognition of the relevance of the fact that: if all of the As are Bs then there might be Bs that are not As. With P1, people see the *weakness* of the path A to C in (7) because there might be Bs that are not As and therefore it is quite possible that no As are Cs.




P2. Recognition of the relevance of the fact that: if All of the Bs are Cs then any As that are Bs are also Cs. With P2, people see the *strength* of the path from A to C in (8) because the As that are Bs must also be Cs.



P3. Recognition of the relevance of the fact that: if As are usually Bs then there are potentially many Bs that are not As. With P3, people see the *weakness* of the path from A to C in (9) because there might be many Bs that are not As and therefore it is again possible that no As are Cs.



P1 – P3 allow people to solve some nonmonotonic reasoning problems with answers that concur with those expected in the AI literature. Thus, for example, consider (1), (2), (5), and (10).



Protocols showed that for problems with structures like these, some people gave the expected answer because they recognized that one path was weaker than another due to P1 - P3. There are four important points about P1 - P3:

- a. People as a whole are not influenced by all of P1 - P3 (Ford & Billington, 2000; Ford, 2004; Ford, in press). Most people are influenced by N1 - N9 to some degree. People who are influenced by P1 - P3 do have lapses where they are sometimes influenced by N1 - N9, to varying degrees.
- b. It is relatively easy to identify people who are likely to be influenced by P1, P2, or P3 and to show that their reasoning in nonmonotonic problems with conflict differs from that of other people. Consider (11) - (13).

(11) Given the following two statements:

All of Jim's friends are Tom's friends.

Tom's friends are usually Fred's friends.

Could it be the case that **none of Jim's friends are Fred's friends?**

(12) Given the following two statements:

Mary's friends are usually Ann's friends.

All of Ann's friends are Sue's friends.

Could it be the case that **none of Mary's friends are Sue's friends?**

(13) Given the following two statements:

Jim's friends are usually Tom's friends.


Tom's friends are usually Fred's friends.

Could it be the case that **none of Jim's friends are Fred's friends?**

The deductively correct answer for (11) and (13) is *Yes* and for (12) it is *No*. P1 yields the answer for (11). P2 yields the answer for (12). P3 yields the answer for (13). Ford (in press; also see Ford, 2004) found that 4 out of 17 undergraduates and 7 out of 10 postgraduates and academics could answer both (11) and (12) correctly. Current studies in our laboratory suggest that understanding both (12) and (13) is easier, with 13 out of 20 undergraduates getting both correct. Also, Dieussaert et al. (2004) found that 11 out of 27 undergraduates could answer both (12) and (13) correctly. Ford (in press, also see Ford 2004) has shown that subjects who understand both (11) and (12) make decisions about nonmonotonic reasoning problems based on the logical strength of the competing arguments, thus favouring the negative path in problems with structures like (1) - the Tweety problem - and favouring a "can't tell" response for problems like (6) - the problem involving contradiction. Those who do not understand (11) and (12) give inconsistent responses, relying on the negative factors identified. Current studies in our laboratory show similar differences between groups of subjects who either do or do not understand

both (12) and (13). Thus, for example, those who understand both (12) and (13) are more likely to see the difference between (6) and (10), more consistently and more strongly favouring the logically stronger, negative, path in (10) than other subjects. Dieussaert et al. (2004) have also shown that subjects who understand both (12) and (13) base their conclusions on differences in the logical strength of arguments and that this reasoning is also influenced by modifier strength.

- c. Although the conclusions to problems such as (1), (2), (5), (6), and (10) based on P1 - P3 concur with those given in the AI community, the basis for the conclusions is different. For the AI community, the preferred conclusions are based on the specificity principle. Analyses of the protocols of the subjects studied by Ford and Billington (2000) and Ford (in press) showed no evidence of a notion of specificity like that considered so important in the AI literature. Rather, the conclusions people gave that concurred with those of the AI community were related to seeing differences in the logical strength of paths: P1 leads to seeing the weakness of $A \rightarrow B \dashrightarrow C$, P2 leads to seeing the strength of $A \dashrightarrow B \rightarrow C$, and P3 leads to seeing the weakness of $A \dashrightarrow B \dashrightarrow C$.
- d. Conclusions based on the specificity principle and those based on P1 - P3 do not always concur. Thus, consider (14) - (16), where i is used to denote an individual who is both an A and a D.



When people use P1 - P3 they seem to assume that all the information given is relevant and that all the relevant information is given and for these problems they will favour the negative paths in (14) and (15) because they are stronger, but will favour “can’t tell” for (16) where the paths are equally strong (Ford & Billington, 2000; Ford, in press; see also Dieussaert et al., 2004; Ford, 2004). In contrast, theories emphasizing the specificity principle will give a “can’t tell” answer for (14) - (16) because the specificity principle does not apply, there being no inheritance path from D to A or from D to B in any of the problems.

We have, then, an example of where at least some “ordinary people” reason in a way that is different from that considered by AI researchers and where new insight into nonmonotonic reasoning is gained which could be

incorporated into artificial systems. Ford (2004) has in fact developed a formal system of nonmonotonic reasoning - known as System LS for logical strength - that goes beyond what people do, but that takes its inspiration from the way that people who understand P1 - P3 make logically justifiable conclusions about nonmonotonic reasoning problems. It is not possible to discuss the system fully here. However, some important points can be noted. System LS is a system of rules that includes the rules of System P (Kraus et al., 1990), which is well-known in the AI community as giving the minimal rules that a reasonable nonmonotonic system must have, but extends them and allows for conclusions at different strength levels (see also Pfeifer & Kleiter, 2005). System LS gives the answers usually captured by the specificity principle. It also gives conclusions preferred in the nonmonotonic logic literature for more complex problems than those considered here (see Ford, 2004). However, it also gives other logically justifiable responses not currently captured in other systems and draws more subtle distinctions between possible conclusions because different levels of strength are recognized. It does this because it is based on the notion of the logical strength of arguments which seems important in logically justifiable human nonmonotonic reasoning. By ignoring the negative things people do, and taking the positive things, a novel, logically justifiable, system has been developed.

Conclusions

It seems that nonmonotonic reasoning should not be considered psychological. It is likely that the apparent ease with which people deal with everyday reasoning is an illusion. They may jump to conclusions easily and retract them if necessary but there may be little justifiable reasoning taking place to resolve conflict. In areas of life where reasoning is crucial but often nonmonotonic, and where mistakes can be devastating, such as medicine, there is more of an acceptance that people do not reason well in everyday life, as noted in the *Journal of the American Medical Association*:

The complex nature of cognition, the vagaries of the physical world, and the inevitable shortages of information processing and schemata ensure that normal humans make multiple errors every day. Slips are most common, since much of our mental functioning is automatic, but the rate of error in knowledge-based processes is higher.

(Leape, 1994, p. 1854).

However, it is still worthwhile studying human nonmonotonic reasoning to help inform artificial systems. Although conclusions drawn by people are often based on flawed assumptions and erroneous notions, at least some people some of the time make judgements that have a consistent, coherent, prin-

cipld basis. By taking the insights of these people, new ideas for artificial nonmonotonic reasoning can be gained and can lead to new systems such as System LS. Further study of human nonmonotonic reasoning should give greater insight into possibilities for reasoning in artificial systems.

References

- Antoniou, G. (1997). *Nonmonotonic reasoning*. MIT Press, Cambridge, Massachusetts.
- Bacchus, F. (1989). A modest, but semantically well founded, inheritance reasoner. *Proceedings of the 11th International Joint Conference on AI* (pp. 1104-1109).
- Benferhat, S., Bonnefon, J. F., & Da Silva Neves, R. (in press). An overview of possibilistic handling of default reasoning, with experimental studies. *Synthese*.
- Brachman, R. J. (1990). The future of knowledge representation. In *Proceedings of the 8th National Conference on Artificial Intelligence AAAI-90* (pp.1082-1092). Menlo Park, CA: AAAI Press.
- Dieussaert, K., Ford, M., & Horsten, L. (2004). Influencing nonmonotonic reasoning by modifier strength manipulation. In K. Forbus, D. Gentier, & T. Regier (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 315-320). Chicago, ILL: Cognitive Science Society, Inc.
- Elio, R. & Pelletier, F. J. (1993). Human benchmarks on AI's benchmarks problems. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society* (pp. 406-411). Hillsdale, NJ.: Lawrence Erlbaum.
- Elio, R. & Pelletier, F. J. (1994). On relevance in non-monotonic reasoning: Some empirical studies. In R. Greiner & D. Subramanian (Eds.), *Relevance: American Association for Artificial Intelligence 1994 Fall Symposium Series* (pp. 64-67). Menlo Park, California: AAAI Press.
- Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. J. (1993). *Human reasoning: The psychology of deduction*. Hillsdale, NJ: Lawrence Erlbaum.
- Ford, M. (1995). Two modes of mental representation and problem solution in syllogistic reasoning. *Cognition*, 54, 1-71.
- Ford, M. (2004). System LS: A three-tiered nonmonotonic reasoning system. *Computational Intelligence*, 20, 89-108.
- Ford, M. (in press). Human nonmonotonic reasoning: the importance of seeing the logical strength of arguments. *Synthese*.
- Ford, M. & Billington, D. (2000). Strategies in human nonmonotonic reasoning. *Computational Intelligence*, 16, 446-468.
- Gabbay, D. (1994). Preface. In D. M. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of logic in artificial intelligence and logic programming: Vol. 3: Nonmonotonic reasoning and uncertain reasoning* (pp. v-xii). Oxford: Clarendon Press.
- Ginsberg, M. (ed.) (1987). *Readings in nonmonotonic reasoning*. Los Altos California: Morgan Kaufmann Publishers.
- Hewson, C. & Vogel, C. (1994). Psychological evidence for assumptions of path-based inheritance reasoning. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 409-414).

- Horty, J. F. (1994). Some direct theories of nonmonotonic reasoning. In D. M. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of logic in artificial intelligence and logic programming, Vol. 3: Nonmonotonic reasoning and uncertain reasoning* (pp. 111-187). Oxford: Clarendon Press.
- Horty, J. F. & Thomason, R. H. (1988). Mixing strict and defeasible inheritance. In *Proceedings of the 7th National Conference on Artificial Intelligence AAAI-88* (pp. 427-432). Menlo Park, California: AAAI Press.
- Johnson-Laird, P. N. & Byrne, R. M. J. (1991). *Deduction*. Hove: Erlbaum.
- Kraus, S., Lehmann, D., & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative models. *Artificial Intelligence*, 44, 167-207.
- Leape, L. L. (1994). Error in medicine. *Journal of the American Medical Association*, 272, 1851-1857.
- Lukaszcwicz, W. (1990). *Non-monotonic reasoning: formalization of commonsense reasoning*. New York: Ellis Horwood.
- Nute, D. (1994). Defeasible logic. In D. M. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming: Vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning* (pp. 353-395). Oxford: Clarendon Press.
- Oaksford, M. & Chater, N. (1998). *Rationality in an uncertain world: Essays on the cognitive science of human reasoning*. Hove, Sussex, UK: Psychology Press.
- Pelletier, F. J. & Elio, R. (1997). What should default reasoning be, by default? *Computational Intelligence*, 13, 165-187.
- Pelletier, F. J. & Elio, R. (in press). The case for psychologism in default and inheritance reasoning. *Synthese*.
- Pfeifer, N. & Kleiter, G. D. (2003). Nonmonotonicity and human probabilistic reasoning. In *Proceedings of the 6th workshop on uncertainty processing*, Hejnice, Czech Republic.
- Pfeifer, N. & Kleiter, G. D. (in press). Coherence and nonmonotonicity in human reasoning. *Synthese*.
- Pfeifer, N. & Kleiter, G. D. (2005). Towards a mental probability logic. *Psychologica Belgica*, 45 (1), 71-99.
- Schlechta, K. (1997). *Nonmonotonic logics: Basic concepts, results, and techniques*, Berlin: Springer-Verlag.
- Schurz, G. (2002). Nonmonotonic reasoning: Ontic, logical and cognitive foundations. *1st Salzburg Workshop on Paradigms of Cognition*, Salzburg, July 2002. (Also to appear in Schurz, G. (in press). Non-monotonic reasoning from an evolution-theoretic perspective: ontic, logical and cognitive foundations. *Synthese*).
- Stein, L. A. (1989). Skeptical inheritance: Computing the intersection of credulous extensions. In *Proceedings of the 11th International Joint Conference on AI* (pp. 1153-1158).
- Touretzky, D. S. (1984). Implicit ordering of defaults in inheritance systems. In *Proceedings of the 5th National Conference on Artificial Intelligence* (pp. 322-325). Menlo Park, CA: AAAI-84.
- Touretzky, D. S., Horty, J. F., & Thomason, R. H. (1987). A clash of intuitions: the current state of nonmonotonic multiple inheritance systems. *Proceedings of 10th International Joint Conference on AI* (pp. 476-482).
- Vogel, C. (1996). Human reasoning with negative defaults. In D. Gabbay & H. J. Ohlbach (Eds.), *Lecture notes in artificial intelligence: Practical reasoning* (pp. 606-621). Berlin: Springer.