Confidence as Higher Order Uncertainty

Pei Wang
Center for Research on Concepts and Cognition, Indiana University
510 North Fess Street, Bloomington, IN 47408
pwang@cogsci.indiana.edu

Abstract

With insufficient knowledge, the conclusions made by a reasoning system are usually uncertain. If the system is open to new knowledge, it also suffers from a higher order uncertainty, because the first order uncertainty evaluations are uncertain themselves — they can be changed by future evidence.

Several approaches have been proposed for handling higher order uncertainty, including the Bayesian approach, higher-order probability, and so on. Though each of them has its advantages, none of them is satisfactory, for various reasons.

A new measurement, confidence, is defined to indicate higher order uncertainty, which is understood as relative stability of first order uncertainty evaluation, and is processed accordingly.

1 Introduction

Non-Axiomatic Reasoning System (NARS for short) is an intelligent reasoning system ([20, 21]). As a reasoning system, it accept knowledge from its environment in a formal language, and answer questions according to its knowledge. As an intelligent system, it works under the assumption of insufficient knowledge and resources . More concretely, the following assumptions are made about its working environment:

1. The system's computing power, as well as its working and storage space, is limited and often in short supply;
2. The tasks that the system has to process, that is, new knowledge and questions, can emerge at any time, and all questions have deadlines attached with them;

3. The system not only can retrieve relevant knowledge and derive sound conclusions from it, but also can make defeasible hypotheses and guesses based on it when no certain conclusion can be drawn;

4. No restriction is imposed on the relationship between old knowledge and new knowledge, as long as they are representable in the system's interface language.

NARS can adapt its behavior according to its experience, that is, to accommodate itself to new knowledge, and to adjust its memory structure and resources distribution to improve its estimated time and space efficiency, under the assumption that future situations will be similar to past situations.

These assumptions are chosen because of their theoretical importance (they can explain many aspects of intelligent behaviors) and their practical usage (many domains have these properties). For a more detailed discussion about the assumptions, see [21].

It follows directly from the assumptions that the system's judgments are usually uncertain, since the input knowledge is not necessarily conflicts-free, and the system need to make plausible inferences when the available knowledge is incomplete for a judgment task.

As a result, for a given question, the system usually cannot find a unique "correct" or "optimal" answer, but a set of uncertain, competing answers. To make a reasonable choice among them, a quantitative measurement for uncertainty is necessary.

Let's assume that there is a well-defined way to measure the *weight of evidence* for a statement (for how such a measurement is formally defined, see [20]). It is natural to judge the uncertainty of the statement by the *frequency* (or *proportion*) of available positive evidence whose weight is w^+, among all relevant evidence whose weight is w, that is, by $\frac{w^+}{w}$.

Since the system is always open to new knowledge, and the system works in an incremental manner, that is, to take knowledge into consideration piece by piece, the *frequence evaluation* itself is uncertain, too — it need to be adjusted according to new knowledge or further consideration.

For example, "Birds can fly" is a statement, then "90% birds can fly" is a *higher order* statement, which expresses the uncertainty (according to the system's available knowledge) of the original statement. Because "90%" is an estimation that is changeable in the future, the higher order statement is uncertain, too.

It is easy to see that the *higher order uncertainty* is also a matter of degree. In the above case, we can easily distinguish the following two situations:

1. 10 birds are observed, and 9 of them can fly; and
2. 10000 birds are observed, and 9000 of them can fly.

Obviously, they lead to different *higher order uncertainty* for "90% birds can fly", and the latter is much more certain than the former in the higher-order sense, though their uncertainty are the same at the first-order, that is, in the two situations "Birds can fly" has the same frequency.

Therefore, NARS do need a measurement about higher order uncertainty of judgments, which indicate *how easily a frequency evaluation can be changed*, so it is related to the concepts of confidence, ignorance, credibility, reliability, stability, sensitivity, susceptibility, and so on.

2 Why to define a new measurement

At the beginning, we may (and I did) expect that there is a ready-made mathematical tool to represent and process this type of uncertainty, since the idea of higher order uncertainty is not novel at all. In the following, let's check several approaches suggested previously, to see whether they can be applied to the situation of NARS.

According to the advocates of the Bayesian paradigm, there is no need to introduce a new measurement, since the information about the higher order uncertainty, often referred as *confidence* or *ignorance*, is a "build-in feature" of a probability distribution function, though the information is implicitly represented there ([3, 14, 18]). Since the higher order uncertainty is actually about the susceptibility of a probability assignment $BEL(E)$ in light of future evidence, they suggest to "to associate partial confidence in $BEL(E)$ with the susceptibility of $BEL(E|c)$ to the various contingencies in C", then confidence can be measured by the "fluctuations in $BEL(E|c)$" or the "narrowness of the distribution of $BEL(E|c)$" ([14]).

What is wrong about this approach, as argued in [19], is the assumption that all re-evaluation of $BEL(E)$, caused by new knowledge c, can be put into the form $BEL(E|c)$, that is, by conditionalization on c.

This assumption is not always valid, since in the above formula c must satisfy the following constraints: (1) c is a binary proposition, (2) c is already in the proposition space upon which $BEL(E)$ is defined, and (3) $BEL(c) > 0$.

Extending Bayes' Theorem to Jeffrey's Rule doesn't solve all the problem, but make another problem of this approach more obvious: the operation is *updating* (by which one judgment is replaced by a competing one, then other judgments are adjusted accordingly), rather than *revision* (by which two competing judgments are combined in a symmetric way) ([4, 19]).

Therefore, *confidence* defined in this way only reflects the stability of a probability (or frequency) assignment to *certain relevant evidence*, and the restrictions upon new knowledge will severely limit the learning ability of the system. Especially, they make the system only open to certain types of new knowledge, therefore inconsistent with the definition of NARS.

Generally speaking, as argued in [19], the higher order uncertainty discussed above cannot be derived from a first order probability distribution, because it is about the background knowledge of the distribution, with is not totally accessible from the distribution function itself. Therefore, we do need a measurement which is specially for this type of uncertainty.

One natural idea is to apply probability theory once again, which leads to the concepts like "probability of probability", "second order probability", "higher order probability", and so on.

This type of approaches have been proposed by several authors ([8, 9, 13]). Though they have the advantage that probability theory provide a solid mathematical foundation, their semantics and utility have been challenged strongly ([12, 14]). In this paper, I only want to argue that they are (at least) inapplicable to the representing and processing of the higher order uncertainty described above, for the following reasons:

1. For a "second order probability" to make sense, it is necessary to assume the existence of an "objective first order probability", that is, the *frequency* of a judgment always has a limit. However, such an assumption is inconsistent with the "insufficient knowledge" assumption accepted by NARS.

2. Even when such an objective probability exist, it is impossible for the system to know how close the current frequency is to it, because such accuracy information is not available to the system. If the second order probability is interpreted as an estimation itself, a "third order probability" will be introduced, $\cdots$ so to cause an infinite regression ([15]).

3. Even if a second order probability can be properly interpreted, it is still not the measurement of confidence or ignorance discussed above.

> Let's say that p and q measure the first and second order of uncertainty of statement S, respectively. When $q = 1$, it means the same thing when interpreted as "second order probability" or "degree of confidence", that is, p is the "true probability" of S in the sense that it will not be changed by future evidence. However, when $q = 0$, the two interpretations are different ([23]). If q is a probability, a 0 means "the probability of 'the probability of S is p' is 0". If q is a measurement of confidence, indicating how sensible p is to future evidence, a 0 means "the system know nothing about the probability of S".

There are other attempts to measure the higher order uncertainty, and keep it different from a probability of probability. For examples, both Shafer's *reliability* and Yager's *credibility* are such measurements, where 0 is interpreted as "The frequency value is unknown", rather than "The frequency value is incorrect". These approaches reduce the higher order uncertainty in a judgment either to the reliability of its information source ([17, 23]), or to its compatibility with higher priority evidence ([23]). Though these two factors do influence the stability of a frequency value, they can hardly explain all the related phenomena. In many situations, it is possible for information provided by the same source to have different stability, and the difference can be detected before the information is compared with background knowledge to check their compatibility. For example, if after observing 10000 birds, you find that 90% birds can fly, you are more confident to say "90% birds can fly" than after observing 9 flying birds among 10 of them, regardless your evaluation about the reliability of your eye. Generally speaking, the principal factor for the higher order uncertainty is the *amount of available evidence*, to which other factors, like the reliability of information sources and compatibility with background knowledge, can be used to make further "discount".

Confidence measurements are also introduced in the study of human judgment and decision making under uncertainty, where it is often related to the accuracy of predictions, that is, for all propositions assigned a given probability q, whether q% of them are really true ([6, 7]). This kind of measurement has its value in psychological studies, but they cannot be used on our current situation, because the "accuracy of probabilistic predictions" and the "stability of probabilistic predictions" don't determine each other. On the other hand, such an interpretation of confidence presumes that every probabilistic prediction will become either true or false, after the related event happens. Such a presumption is not shared by NARS.

3 How is confidence defined in NARS

For the reasons discussed in the previous section, a new measurement of confidence is introduced in NARS.

As discussed at the beginning of the paper, the higher order uncertainty appears as the result of insufficient knowledge. For the same reason, it doesn't make sense to talk about an "objective" or "correct" frequency, and to use its relation to the current frequency as a measurement of higher order uncertainty. The current frequency value is uncertain, not because it is an estimation of an "objective value", but because it will be influenced by future evidence, and its stability is a matter of degree.

How to measure the stability (or its contrary, susceptibility) of a frequency value f? A natural idea is by "how much it will be changed by future evidence". Because NARS is always open to new evidence, and new evidence may conflict with current belief, f can be anywhere in $[0, 1]$ in the infinite future, no matter what its current value is. So no frequency is stable in the absolute sense.

However, what we are interested in is the *relative stability* of different frequency assignments, that is, given the same amount of new evidence, how much each of them will change.

Let's say (as defined in the first section) that the current frequency of a statement is $f = \frac{w^+}{w}$, $0 \leq w^+ \leq w$, where w^+ and w are the weight of positive and total relevant evidence, respectively. If in the *near future* some new evidence is available, and its weight is k ($k > 0$), then where f will be?

Obviously, if the new evidence is completely negative, f will be $\frac{w^+}{w+k}$; if the new evidence is completely positive, f will be $\frac{w^+ + k}{w+k}$. Therefore, no matter what content that amount of evidence has, the frequency will stay in the interval $[\frac{w^+}{w+k}, \frac{w^+ + k}{w+k}]$. The *width* of the interval, $\frac{k}{w+k}$, provides a good measurement for the ignorance (or susceptibility) of the judgment, and its complement (to 1), $\frac{w}{w+k}$, provides a good measurement for the confidence (or stability) of the judgment.

To use the width of an interval to represent ignorance is not a novel idea, and it has been accepted by the "probability interval" approaches ([2, 10, 11]) and Dempster-Shafer theory ([16]). What makes NARS different from the other approaches is the definition of the interval: here it is the interval where *the frequency will be in the near future* (see [22] for why the other interval definitions cannot be used in NARS).

Let's take $k = 2$ (see [20] for an further explanation of k), if 10 birds are observed and 9 of them can fly, then "Birds can fly" has a frequency 0.9000 and a confidence 0.8333; if 10000 birds are observed and 9000 of them can fly, then "Birds can fly" has a frequency 0.9000 and a confidence 0.9998.

Let c be $\frac{w}{w+k}$ is consistent with our previous discussions about confidence:

1. $c = 0$ is identical with $w = 0$, that is, no evidence, maximum ignorance, minimum confidence. The future frequency will be completely determined by new evidence.

2. $c = 1$ is identical with $w \to \infty$, that is, infinite evidence, minimum ignorance, maximum confidence. The current frequency will no longer be influenced by new evidence. Such a situation cannot be reached by the accumulation of evidence, but can be used to represent definitions and conventions in the system.

3. c increases monotonically with w, that corresponds to the psychological phenomenon that confidence "increase as a function of of the amount of information available" ([1, 6]).

4. Though c can be represented as a *ratio*, that is, *the weight of evidence the system has at current* to *the weight of evidence the system will have in the near future*, it is *not* a probability, since it doesn't follow the axioms of probability theory.

5. The higher a judgment's confidence is, the harder the judgment's frequency can be changed by future evidence. However, it doesn't mean that the judgment is "more accurate" ([6, 7]) in an objective sense.

6. The *frequency* and *confidence* of a judgment are *independent* to each other, that is, from the value of one, the other's value cannot be determined, or even estimated or bounded.

The $< f, c >$ pair is referred to as the truth value of a judgment in NARS. According to the previous discussion, such a truth value cannot have a model-theoretic semantics, that is, it doesn't tell us to what extent the statement matches "state of affairs". However, it can tell us to what extent the statement is supported by available knowledge. This is what we can get under the assumption of insufficient knowledge and resources.

Defined in this way, there is no "third-order uncertainty" to worry about. The "stability" of a confidence value can be derived from the confidence value itself. Because the current confidence is $c = \frac{w}{w+k}$, in the near future, with the coming of new evidence whose weight is k, the new confidence will be $\frac{w+k}{w+2k}$, that is, $\frac{1}{2-c}$. So we don't need another measurement, and there is no infinite regression.

The last question is: is this kind of information available for the system? Even Bayesian network and fuzzy logic, both require the users to assign *a single number* to each proposition, meet the objections of "nowhere to get the numbers", how can we expect the users to provide *a pair of numbers* for each judgment?

To me, the hardness of value assignments comes mainly from the unclear interpretation about *what are measured by these values.* NARS attempts to be user friendly by unifying different uncertainty representations, so to make the users easier to understand them. Users can even mix different forms of truth value, in terms of *weight of evidence, frequency, confidence, ignorance, frequency interval,* and so on, in the knowledge they provided. NARS also accept truth values represented as *a single number* (by taking accuracy into consideration) or a *linguistic variable* (by translate it into an frequency interval). A detailed description on this issue can be found in [22].

4 How is confidence processed in NARS

After given a definition and interpretation of confidence, let's see how it is processed in by the various inference rules in NARS. In the following discussion, we'll concentrate on confidence. For a more complete description of NARS, see [20, 21].

4.1 Negation

If the truth value assigned to a statement is $<f,\, c>$, then what truth value should be assigned to the negation of the statement?

According to above discussion, we know that $f = \frac{w^+}{w}$, and $c = \frac{w}{w+k}$. By definition, the positive evidence of a statement is negative evidence for its negation, and *vice versa.* Therefore, the truth value for the negated statement is $< 1 - f, c >$. Here we can see again that f is a probability function, but c is not.

For example, if "Birds can fly" has a frequency 0.9000 and a confidence 0.8333, that means its weight of positive evidence is 9, and its weight of total evidence is 10. Therefore, the positive evidence for "Birds cannot fly" has a weight 1, and the total evidence is unchanged. Consequently, it get a frequency 0.1000 and a confidence 0.8333.

4.2 Expectation

How to estimate the future frequency from the past frequency and confidence? From probability theory we know that with a large sample space, we can simply use the past frequency as our expectation e, that correspond to the case when c is close to 1 in NARS. For a small sample size, say w^+ successes in w tests, a popular formula used to estimate the probability of success in the next test is *Laplace's low of succession*: $e = \frac{w^+ + 1}{w+2}$.

NARS uses a generalization of this rule, $e = \frac{w^+ + \frac{k}{2}}{w+k}$, where k is the constant mentioned above. When expressed as a function of the truth value, we get $e = c\,(f - \frac{1}{2}) + \frac{1}{2}$. Intuitively, f is "squashed" by the factor c to the "no preference point" $\frac{1}{2}$ to become e. As results, we have

1. When $c = 0$, $e = 0$. That is, with null evidence, the system has no preference on whether a statement can be verified by future evidence.

2. When $c = 1$, $e = f$. That is, with complete evidence, the expectation equals the frequency. As discussed previously, this corresponds to the situation that f is not come from empirical evidence, but from definition or convention.

3. In all other cases, e is always "more conservative" (closer to the "no preference" point) than f. This conservatism can be explained by the consideration that past experience may be different from future experience. The smaller c is, the more conservative the result is. Similar conservatism has been observed in human information processing behaviors ([5]).

4.3 Revision and updating

In NARS, *revision* indicates the process by which evidence from different sources are combined.

For example, assuming the system's previous truth value for "Birds can fly" is $< 0.9000, 0.8333 >$ (we know that it corresponds to "10 birds, 9 can fly"), now a piece of new knowledge comes, which is "birds can fly" $<0.7500, 0.6667>$ (so it corresponds to "4 birds, 3 can fly"). If the system can determine that no evidence is repeatedly counted in the two sources (see [20] for how this is defined and checked), then the truth value of the revised judgment should be $<0.8571, 0.875>$ (corresponding to "14 birds, 12 can fly").

The revision function, represented as from a pair of truth values of the premises to that of the conclusion, has the following form:

$$f = \frac{w_1 f_1 + w_2 f_2}{w_1 + w_2}$$

$$c = \frac{w_1 + w_2}{w_1 + w_2 + k}$$

where $w_i = k\frac{c_i}{1-c_i}$ is the weight of total evidence of judgment i ($i = 1, 2$).

The revision function has the following properties:

1. The order of the premises doesn't matter.

2. As a weighted average of f_1 and f_2, f is usually a "compromise" of them, and is closer to the one that is supported by more evidence.

3. c is smaller than neither c_1 nor c_2, that is, the conclusion is supported by no less evidence than a premise.

4. If $c_1 = 0$, then $f = f_2$ and $c = c_2$, that is, a judgment supported by null evidence cannot revise other judgment.

5. If $c_1 = 1$ and $c_2 < 1$, then $f = f_1$ and $c = c_1$, that is, a definition (supported by complete evidence) cannot be modified by empirical evidence.

What will happen if the evidence of the two premises are "correlated", that is, some evidence are repeatedly counted in them? In such a case, NARS applies the *updating* rule, to pick up the premise with a higher *confidence*, since it is supported by more evidence. Such a confidence-based updating is different from the updating in the Bayesian approach, where *new* evidence always suppress *old* evidence ([4, 19]).

4.4 Syllogisms

The syllogisms in NARS are rules for deduction, abduction, and induction. These rules also include functions calculating the truth value of the conclusions from those of the premises. The concrete form of the rules can be found in [20], which is beyond the scope of this paper. In the following, I only mention two facts about these rules that is related to confidence:

1. The confidence of a conclusion is not larger than the confidence of either premise, that is, confidence "declines" in syllogistic inference.

2. Confidence declines much slower in deduction than in induction and abduction. In deduction, if both premises have a confidence value 1, the conclusion may also have a confidence value 1 (so it is a derived definition or convention). In induction and abduction, as a contrary, the confidence of the conclusion has a upper bound which is far less than 1. So, by saying "induction and abduction are more uncertain when compared with deduction", what is referred to is not the "first-order uncertainty" f, but the "higher-order uncertainty" c.

5 Summary

With insufficient knowledge, the conclusions made by a reasoning system are usually uncertain. If the system is also open to new knowledge, there is a higher order uncertainty, which indicating the stability of first order uncertainty evaluations.

Several approaches have been proposed for handling higher order uncertainty. Though each of them has its suitable application domain, they are not appropriate for the uncertainty described above, for various reasons.

Confidence is defined in NARS as a measurement of higher order uncertainty, which is understood as relative stability of first order uncertainty evaluation (frequency), and is processed according to such an interpretation. It is also defined in such a way that closely related to other uncertainty measurements.

Since NARS is designed as a general purpose intelligent reasoning system, the confidence measurement is domain-independent.

Like all other approaches, the NARS approach for uncertainty representation is based on certain assumptions about the environment. For NARS, the fundamental assumptions are: the system's knowledge and resources are usually insufficient, and the environment is relatively stable. Concretely, its confidence measurement and processing are based on the availability of an additive *weight of evidence* function. This function is not really used to evaluate each piece of evidence, but to provide a semantic interpretation for truth values.

NARS is not necessarily better than the competing approaches in all kinds of environment, but (hopefully) is better in the environment described at the beginning of the paper, which has special theoretical and practical interests from the view point of artificial intelligence and cognitive science.

Acknowledgment

This work is supported by a research assistantship from the Center for Research on Concepts and Cognition, Indiana University.

References

[1] N. Anderson. A cognitive theory of judgment and decision. In B. Brehmer, H. Jungermann, P. Lourens, and G. Sevón, editors, *New Directions in Research on Decision Making*, pages 63–108. Elsevier Science Publishers, Amsterdam, 1986.

[2] P. Bonissone. Summarizing and propagating uncertain information with triangular norms. *International Journal of Approximate Reasoning*, 1:71–101, 1987.

[3] P. Cheeseman. In defense of probability. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 1002–1009, 1985.

[4] D. Dubois and H. Prade. Updating with belief functions, ordinal conditional functions and possibility measures. In P. Bonissone, M. Henrion, L. Kanal, and J. Lemmer, editors, *Uncertainty in Artificial Intelligence 6*, pages 311–329. North-Holland, Amsterdam, 1991.

[5] W. Edwards. Conservatism in human information processing. In B. Kleinmuntz, editor, *Formal Representation of Human Judgment*, pages 17–52. Wiley, New York, 1968.

[6] H. Einhorn and R. Hogarth. Confidence in judgment: persistence of illusion of validity. *Psychological Review*, 35:395–416, 1978.

[7] B. Fischhoff, P. Solvic, and S. Lichtenstein. Knowing with certainty: the appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance*, 3:552–564, 1977.

[8] R. Fung and C. Chong. Metaprobability and Dempster-Shafer in evidential reasoning. In L. Kanal and J. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 295–302. North-Holland, Amsterdam, 1986.

[9] H. Gaifman. A theory of higher order probabilities. In J. Halpern, editor, *Theoretical Aspects of Reasoning about Knowledge*, pages 275–292. Morgan Kaufmann, Los Altos, California, 1986.

[10] B. Grosof. An inequality paradigm for probabilistic knowledge. In L. Kanal and J. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 259–275. North-Holland, Amsterdam, 1986.

[11] H. Kyburg. Bayesian and non-Bayesian evidential updating. *Artificial Intelligence*, 31:271–293, 1987.

[12] H. Kyburg. Higher order probabilities. In L. Kanal, T. Levitt, and J. Lemmer, editors, *Uncertainty in Artificial Intelligence 3*, pages 15–22. North-Holland, Amsterdam, 1989.

[13] G. Paaß. Second order probabilities for uncertain and conflicting evidence. In P. Bonissone, M. Henrion, L. Kanal, and J. Lemmer, editors, *Uncertainty in Artificial Intelligence 6*, pages 447–456. North-Holland, Amsterdam, 1991.

[14] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann Publishers, San Mateo, California, 1988.

[15] L. Savage. *The Foundations of Statistics.* Wiley, New York, 1954.

[16] G. Shafer. *A Mathematical Theory of Evidence.* Princeton University Press, Princeton, New Jersey, 1976.

[17] G. Shafer. Perspectives on the theory and practice of belief functions. *International Journal of Approximate Reasoning*, 4:323–362, 1990.

[18] D. Spiegelhalter. A statistical view of uncertainty in expert systems. In W. Gale, editor, *Artificial Intelligence and Statistics*, pages 17–56. Addison Wesley, Reading, 1986.

[19] P. Wang. Belief revision in probability theory. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, pages 519–526. Morgan Kaufmann Publishers, San Mateo, California, 1993. Available via WWW at *http://www.cogsci.indiana.edu/farg/pwang_papers.html*.

[20] P. Wang. From inheritance relation to non-axiomatic logic. Technical Report 84, Center for Research on Concepts and Cognition, Indiana University, Bloomington, Indiana, 1993. Available via WWW at *http://www.cogsci.indiana.edu/farg/pwang_papers.html*. A revised version is going to appear in the *International Journal of Approximate Reasoning*.

[21] P. Wang. Non-axiomatic reasoning system (version 2.2). Technical Report 75, Center for Research on Concepts and Cognition, Indiana University, Bloomington, Indiana, 1993. Available via WWW at *http://www.cogsci.indiana.edu/farg/pwang_papers.html*.

[22] P. Wang. An unified treatment of uncertainties. Technical Report 87, Center for Research on Concepts and Cognition, Indiana University, Bloomington, Indiana, 1993. Available via WWW at *http://www.cogsci.indiana.edu/farg/pwang_papers.html*.

[23] R. Yager. Credibility discounting in the theory of approximate reasoning. In P. Bonissone, M. Henrion, L. Kanal, and J. Lemmer, editors, *Uncertainty in Artificial Intelligence 6*, pages 299–310. North-Holland, Amsterdam, 1991.