Stereotypical Reasoning: Logical Properties

DANIEL LEHMANN, *Institute of Computer Science, Hebrew University, Jerusalem 91904, Israel. E-mail: lehmann@cs.huji.ac.il*

Abstract

Stereotypical reasoning assumes that the situation at hand is one of a kind and that it enjoys the properties generally associated with that kind of situation. It is one of the most basic forms of nonmonotonic reasoning. A formal model for stereotypical reasoning is proposed and the logical properties of this form of reasoning are studied. Stereotypical reasoning is shown to be cumulative under weak assumptions.[1]

Keywords: prototypical reasoning, stereotypical reasoning, nonmonotonic consequence relations

1 Introduction

Common sense reasoning in AI requires drawing inferences in a bolder, more adventurous way, than mathematical reasoning. Many different formalisms that implement some form of bold reasoning have been proposed, implemented, used to build artificial systems. Almost no work has been done comparing those formalisms with the way natural intelligence deals with those tasks. Minsky [5, 6] represents probably one of the only efforts to model reasoning performed by natural intelligence.

During the last decades, philosophers, linguists, sociologists have revolutionized the way we understand the human mind. Putnam [7] has criticized the classical philosophical assumptions about meaning, and claimed that stereotypes are a necessary component of the meaning of terms. Rosch [8, 9, 10] has put in evidence the essential function of categorization in achieving intelligence and the intricate ways in which we use it. Categorization is the process in which we relate a specific object or situation to the kind we shall think it a member of. She showed that many of our categories have prototypes, i.e., best examples. Lakoff [3] resumes and expands much of this line of research.

The purpose of this work is to begin the study of inferencing in a mind that uses categories as described above. A model of the simplest kind of inferencing using stereotypes will be given and the formal properties of the inferencing process will be studied. The formal properties of inferencing that are of interest have been singled out by [1, 2, 4].

[1] Full version of an invited paper presented at the *4th Workshop on Logic, Language, Information and Computation* (*WoLLIC'97*), held in Fortaleza (Ceará), Brazil, August 19–22 1997, with scientific sponsorship by IGPL, FoLLI and ASL, and organised by Univ. Federal do Ceará (UFC) and Univ. Federal de Pernambuco (UFPE).

2 Stereotypical reasoning

In this work, stereotypical reasoning is used to denote what is probably the simplest form of natural nonmonotonic reasoning. The present use of the word stereotype is very closely related to Putnam's stereotypes. He claimed that stereotypes are a necessary part of the meaning of words denoting a *natural kind*. Here, stereotypes are assumed not only for natural kinds but for any state of information. This could be understood as the assumption that stereotypes are part of the meaning of any sentence, but the philosophical aspects of this assumption are not discussed in this paper. The point of this paper is the study of how stereotypes are used in the inferencing process, and the formal properties of the inferencing resulting from the use of stereotypes. The use of stereotypes has not been discussed by Putnam.

What is called here stereotypical reasoning is very closely related to the use of what Rosch calls prototypical categories. Prototypical alludes, though, to a richer structure than stereotypical and this is the reason the latter term has been preferred. In ordinary parlance stereotypes are considered to be typically wildly inaccurate and an impediment to intelligent thinking. This reputation should not hide the fact that the use of stereotypes is a fundamental tool, probably the central tool, in achieving intelligence. Hence, the importance of its study. Nevertheless, the negative connotation attached to the word *stereotype* should remind us we are studying a limited form of reasoning, certainly not capable of exhibiting all forms of intelligence.

Here is an example of what I will call stereotypical reasoning. The choice of the *tiger* stereotype follows Putnam. If Benjamin tells you that during his trip in India, hiking in the jungle, he saw a tiger, you will assume he saw a large, frightening animal, yellow with black stripes. Note that not all tigers are such. Some tigers are small, dead, or albino. You have been using the stereotype that says that tigers are big, dangerous and yellow with black stripes. The use of this stereotype may be a mistake: the end of the story may reveal this was an albino tiger, but, typically, the use of the stereotype is precisely what enables efficient communication, since Benjamin knows you have this stereotype (as he has) he assumes you will draw the corresponding conclusions and he intends you to draw those conclusions.

This simple example already suggests a number of questions, most of them will not be touched upon in this paper. What is the nature, or the structure of the stereotype *tiger*? Is it just a conjunction (or some other composition) of properties? In this work, we shall assume that yes, a stereotype is a set of possible states of affairs, but the proper treatment of prototypical categories in general may need a more sophisticated structure. Note, though, that we are not assuming (the classical view, attacked by Lakoff and others) that categories are sets (of models), far from that. We are only assuming that stereotypes are sets of models. In fact, the way stereotypes are used makes them function very much like the graded or radial categories of Lakoff.

How are such stereotypes acquired? This is certainly both deeply rooted in our physiology and a social process.

Why is this the *right* stereotype for tiger? Why is it any better than some other? Here, an analysis of rationality and utility is certainly needed.

What made you apply the *tiger* stereotype to the little story above and not, for example, the *jungle* stereotype that says that, in the jungle, everything is dark green, or the *India* stereotype, whatever this is for you? In this paper this choice will be modeled by some distance between the information at hand and the stereotype. We shall not be able to explain why a specific distance is used. Such an explanation would certainly be based on utility considerations.

3 A formal model of stereotypical reasoning

Informally, starting from information about the situation at hand, one chooses the best stereotype to fit the information and uses both the original information and the stereotypical information to draw conclusions.

Formally, we assume W is the set of all possible states of affairs (i.e., models or situations). We assume a collection (not necessarily finite, but it may well be finite) of stereotypes, S_i. Notice we use lower indexes to identify the stereotypes. Each stereotype is a subset of W, the set of situations in which the stereotype holds. For example the *tiger* stereotype, S_{tiger} could be the set of all models in which tigers are frightening live animals, yellow with black stripes.

The user has some information, i.e., facts, about the situation at hand. This information is modeled by a subset F of W: the set of all situations compatible with the information at hand. On the basis of F, the reasoner picks up one of the stereotypes: S^F, the stereotype most appropriate to F, in a way that will be discussed later. Notice we use here an upper index to denote the stereotype that best fits some information F. The reasoner will then conclude that the actual state of affairs is one of the members of the intersection $F' \stackrel{\text{def}}{=} F \cap S^F$. The nonmonotonicity of the reasoning stems from this jump from F to the subset F'. Clearly, we do expect the set F' to be non-empty, assuming F is non-empty, since we want to avoid jumping to contradictory conclusions. It will be the task of the function that defines the best stereotype to pick a stereotype that has a non-empty intersection with the information F at hand. In many cases the facts F are given by a formula α that is known to be true. In this case F is the set of all models that satisfy α. We shall identify the formula α and the set of its models and write S^α for the stereotype most appropriate for the sets of models of α. A formula β is then nonmonotonically deduced from α iff it is satisfied by all elements of F', that is iff any model m in the set S^α that satisfies α satisfies β: $\alpha \mid\!\sim \beta$ iff $\forall m \in S^\alpha$, $m \models \alpha$ implies $m \models \beta$.

Syntactically, this may be described as taking for $\mathcal{C}(X)$, the set of nonmonotonic consequences of a set X of formulas, the set $\mathcal{C}n(X, g(X))$, of all formulas that logically follow from the set $X \cup g(X)$, where $g(X)$ is the set of formulas that hold in all models of the stereotype that best fits X.

Our analysis of stereotypical reasoning will use the simplistic model just described, since it is good enough for the purpose of this paper. If one thinks of first-order languages and models, one may want to refine this model and associate a stereotype with each one of the objects of the structure: e.g., if our story refers to two tigers about which one has different information, one will perhaps use different stereotypes for each of the tigers: a mother tiger stereotype and a pup tiger stereotype for example.

Before we analyze some consequences of this model, let us point out some of its basic limitations. It is assumed that the conclusions from facts F are drawn by identifying a *unique* stereotype most appropriate for F. One may ask whether this should be the case. Instead of picking up a single stereotype, perhaps one should consider the set of all most appropriate stereotypes and use them all, i.e. their intersection. Indeed the results of next section would hold also in this more general model, but the uniqueness assumption will be needed later. Intuitively it seems to me that we do pick up a unique stereotype, sometimes made up of different stereotypes, but that this composition is almost never the simple juxtaposition, i.e., conjunction of stereotypes. The main reason probably is that such conjunctions are very often empty and we certainly want to avoid drawing inconsistent conclusions from consistent facts. Consider, for example, our tiger above. We did not use both the *tiger* and the *jungle* stereotypes because they clash about the color of the tiger: yellow and black vs. dark green. It may be the case that a very smart reasoner will use a *tiger in the jungle* stereotype, that implies the tiger is barely visible, but this stereotype, though including, somehow, both the *tiger* and the *jungle* stereotypes, cannot be reduced to their conjunction. To avoid premature commitment to a theory of the formation of compound stereotypes, this paper will just assume any set F of facts is associated with a unique stereotype.

4 First consequences of the model

As general as it is, the model presented above has some important consequences for the formal properties of the process of nonmonotonic deduction it defines, i.e., of the consequence relation $\mid\!\sim$.

First, since what is defined by facts F is a set of models, F', the set of nonmonotonic consequences of F, i.e., the set of formulas that hold in all elements of F' is a logical theory, i.e., closed under logical consequence. In other terms, the relation $\mid\!\sim$ satisfies the rules of Right Weakening and And of [2]. Secondly, since F' is a subset of F, any formula that is logically implied by F holds in all elements of F', or, the relation $\mid\!\sim$ satisfies the Reflexivity of [2]. Lastly, since the information at hand is represented, semantically, by a set of models, the relation $\mid\!\sim$ satisfies Left Logical Equivalence.

5 Further assumptions

The main purpose of this work is to consider whether other, more sophisticated, logical properties may be expected from stereotypical reasoning. It is clear that, in the very general model described above, without further assumptions, nothing more can be expected: given any relation $\mid\!\sim$ satisfying Left Logical Equivalence, Reflexivity, Right Weakening and And, one may define, for any formula α, the stereotype S^α to be the set of all models that satisfy all the formulas β such that $\alpha \mid\!\sim \beta$. Since the relation $\mid\!\sim$ is reflexive, all models of S^α satisfy α and, for any α, $S^\alpha \subseteq F^\alpha$. Therefore $F^\alpha \cap S^\alpha = S^\alpha$ and the nonmonotonic consequence relation defined by the model is exactly $\mid\!\sim$.

Our goal is to find some additional, reasonable, assumptions about the set of stereotypes or the way the best stereotype for a set F is chosen that will have interesting consequences on the nonmonotonic logic defined. In fact, the set of stereotypes and its structure does not seem to play an important role here and we shall concentrate on the choice of the best stereotype.

Notice that the mapping from a set F to its best stereotype S^F may be very wild. We do not expect, for example, that $F' \subseteq F$ should imply $S^{F'} \subseteq S^F$. It may well be the case that *robins* is the best stereotype for birds, but the best stereotype for antarctic birds is, for lack perhaps of knowledge of a better one, *vertebrates*.

It is extremely helpful to consider the process of associating to the set F the stereotype S^F as based on some notion of distance between *information sets* and *stereotypes*: the best stereotype for F is the stereotype closest to F:

$$d(F, S^F) \leq d(F, S), \text{ for every stereotype } S. \tag{5.1}$$

Notice that this notion of distance is a bit unusual, since it is defined only from information sets to stereotypes. We shall never use the notion of the distance from a stereotype to an information set. The assumption that our choice is based on some notion of a distance does not limit the generality of our model, since one may always find a suitable distance to fit any choice of best stereotype. The interest of this assumption is that it suggests some natural additional assumptions on the properties of this distance. Those assumptions will be related to logical properties of the nonmonotonic deduction.

Let us suppose D is a partially ordered set (of distances) and that there is a function d that associates an element of D $d(F, S)$ with every set of models F (in fact every set of models that could appear as an information set would be enough) and every stereotype S. A first assumption, already described above, is that this distance always enables us to pick a unique best stereotype.

- *Assumption zero*: for any given information set F there exists a *unique* stereotype S^F such that $d(F, S^F) \leq d(F, S)$ for any stereotype S.

A number of examples of models and choice functions will be described. They may not be very intuitively appealing, but their purpose is to help the reader understand our definitions and prove the consistency of the assumptions that shall be made below. In all examples the set D of distances is taken to be the set of integers, eventually with ∞ added.

EXAMPLE 5.1
There is one stereotype only: S_0 and $S_0 = W$. The exact definition of the distance is irrelevant. Assumption zero is satisfied trivially and, for any F, $S^F = S_0 = W$. Clearly, for any F, $F' = F \cap S^F = F$, and therefore F' is non-empty if F is. The non-monotonic logic defined happens to be monotonic and to be the classical one: $\alpha \mid\!\sim \beta$ iff $\alpha \vdash \beta$.

EXAMPLE 5.2
Assume the set W is finite. Every set $S \subseteq W$ is a stereotype and

$$d(F, S) = \mid S - F \mid - \mid S \cap F \mid,$$

where $\mid A \mid$ indicates the cardinality of the set A. Since $d(F, F) = - \mid F \mid \leq d(F, S)$ for any S, we see that, for any F, $S^F = F$, and therefore assumption zero is satisfied, $F' = F$ and the logic defined is the classical one as in Example 5.1.

EXAMPLE 5.3
Assume the set W is the set of natural numbers. Stereotypes are singletons of W. Distances are defined in the following way: if $n \in F$, $d(F, \{n\}) = n$ and if $n \notin F$, $d(F, \{n\}) = \infty$. Clearly S^F is the singleton that contains the minimal element of F, $\min(F)$ and assumption zero is satisfied. Note also that $F' = \min(F)$ is non-empty if F is non-empty. The model boils down to considering that world m is more probable than world n iff $m < n$. The logic defined results in, given a set of possibilities F, jumping to the conclusion that the most probable one must obtain. This provides a highly nonmonotonic consequence relation.

The next example presents a simple, but natural, family of models.

EXAMPLE 5.4
Assume W is finite and the set of (non-empty) stereotypes S_i, $i = 0, \ldots, k-1$ provides a partition of W, i.e., $\bigcup_{i \in k} S_i = W$ and $S_i \cap S_j = \emptyset$, for any $i \neq j$. Given a set F, we associate with it the stereotype S_j which *covers* F best, i.e., for which the size of the set $S - F$ is minimal. In case this criterion does not define a unique stereotype, choose the stereotype with smallest index. Formally we may define the distance by: $d(F, S_i) = \mid S_i - F \mid + \frac{i}{k}$. The consequence relation defined is nonmonotonic.

After these examples, let us consider interesting properties of the distance d. Since F and S are both sets of models (subsets of W) we may, without loss of generality, assume that $d(F, S) = e(F \cap S, S - F, F - S)$. Three additional assumptions concerning the way the function e depends on each of its three arguments are now natural.

- *Assumption one*: the function e is anti-monotone in its first argument. I mean that if $A \subseteq A'$, then $e(A', B, C) \leq e(A, B, C)$. The relation $\subseteq$ is the subset relation. This assumption is very natural: $d(F, S)$ measures the closeness of F and S: the more they have in common, the closer they are. In most cases we expect that the best stereotype for F should be consistent with F, i.e., have an non-empty intersection with F. If this is the case, our assumption is only slightly stronger: all other things being equal, the best stereotype for F has the largest intersection with F. The set $F \cap S^F$ represents the nonmonotonic consequences of F; we prefer weaker consequences, therefore we prefer to take the set $F \cap S^F$ as large as possible.
- *Assumption two*: the function e depends monotonically on its second argument. Here I mean that if $B \subseteq B'$, then $e(A, B, C) \leq e(A, B', C)$. The second argument, $B = S - F$ measures the set of models compatible with the stereotype but excluded by the information. Notice that the stereotype may be vague, i.e., contain a large number of elements: for example the *bird* stereotype may include birds of many colors, and the information at hand may exclude a lot of those elements: for example we may know the bird we are discussing is yellow. The more such elements are excluded by the information at hand, the less suitable is the stereotype: if too many such elements are excluded a more specific stereotype may be more suitable. In our example, a *yellow bird* stereotype, if there is one such stereotype, should be preferred.
- *Assumption three*: the function e does not depend on its third argument. It seems easy to justify that the function e should depend monotonically on its third argument, by an argument very similar to that used for justifying assumption two. It is perhaps a little less obviously natural that e should not depend at all on its third argument. But, notice that the set $F - S$ is a measure of the strength of our nonmonotonic inference: the larger it is the more nonmonotonic consequences we get in addition to the monotonic ones. The argument just above is to the effect we should not get too many such inferences, but we are certainly interested in getting such nonmonotonic consequences, and should not try to minimize them. Our assumption is that how much nonmonotonicity we get should not be a criterion in choosing the best stereotype.

Assumptions one to three may be summarized by the following: for any F, F' and any stereotypes S, S', if

$$F' \cap S' \subseteq F \cap S \text{ and } S - F \subseteq S' - F' \text{ then } d(F,S) \leq d(F',S'). \quad (5.2)$$

PROOF.

$$d(F,S) = e(F \cap S, S - F, F - S) \leq e(F' \cap S', S' - F', F' - S') = d(F',S')$$

∎

One may notice that Equation 5.2 implies that $d(F,S) = d(F \cap S, S)$. Let us consider the examples above again. In Example 5.1, we may define the distance d to be constant, for example $d(F,S) = 0$. Equation 5.2 is obviously satisfied. In Example 5.2 also, Equation 5.2 is obviously satisfied. In Example 5.3, let us check that Equation 5.2 is satisfied. Since stereotypes are singletons, $F' \cap S' \subseteq F \cap S$ implies that either $F' \cap S' = \emptyset$, or $S' = S = F' \cap S' = F \cap S$. In the first case $d(F',S') = \infty$ and the result holds. In the second case, if $S = \{n\}$, $d(F,S) = n = d(F',S')$. For Example 5.4, if $F' \cap S' \subseteq F \cap S$ then, either $F' \cap S' = \emptyset$, or $S' = S$. If $S - F \subseteq S' - F'$, then either $S - F = \emptyset$ or $S' = S$. If $S' = S = S_i$,

$$d(F,S) \ = \mid S - F \mid + \frac{i}{k} \ \leq \mid S' - F' \mid + \frac{i}{k} \ = \ d(F',S').$$

If $S' = S_j \neq S = S_i$, $F' \cap S' = \emptyset$ and $S - F = \emptyset$,

$$d(F,S) \ = \ \frac{i}{k} \ < \ 1 \ \leq \mid S' \mid \leq \mid S' \mid + \frac{j}{k}.$$

Equation 5.2 is satisfied.

In the sequel we shall assume, sometimes without recalling this explicitly, that the distance d satisfies Equation 5.2.

Our main result is that stereotypical reasoning yielded by a distance that satisfies the four assumptions above: i.e., uniqueness of the closest stereotype, antimonotonicity of the distance $d(F,S)$ in $F \cap S$, monotonicity in $S - F$ and independence from $F - S$, is cumulative [2]. The main result is therefore the following.

THEOREM 5.5
If $F \cap S^F \subseteq F' \subseteq F$, then $S^{F'} = S^F$.

PROOF. Assume $F \cap S^F \subseteq F' \subseteq F$. We must show that, for any stereotype S, we have $d(F', S^F) \leq d(F', S)$. First, since $F \cap S^F \subseteq F'$, we have both $F \cap S^F \subseteq F' \cap S^F$ and $S^F - F' \subseteq S^F - F$, therefore, by Equation 5.2, we have $d(F', S^F) \leq d(F, S^F)$. By Equation 5.1, for any stereotype S, $d(F, S^F) \leq d(F,S)$ and therefore, for any S, $d(F', S^F) \leq d(F,S)$. Using, now, $F' \subseteq F$, we see that $F' \cap S \subseteq F \cap S$ and $S - F \subseteq S - F'$. By Equation 5.2, then, $d(F,S) \leq d(F',S)$, and $d(F', S^F) \leq d(F',S)$, for any stereotype S. ∎

COROLLARY 5.6
The nonmonotonic consequence relation $\mathrel{|\!\sim}$ defined by stereotypical reasoning yielded by a distance that satisfies Equation 5.2 satisfies Cut and Cautious Monotonicity and is therefore cumulative.

PROOF. Suppose $\alpha \mid\!\sim \beta$. Let F be the set of models of α and F' be the set of models of $\alpha \wedge \beta$. The assumption $\alpha \mid\!\sim \beta$ means that all elements of $F \cap S^F$ satisfy β, i.e., $F \cap S^F \subseteq F'$. But clearly $F' \subseteq F$. By Theorem 5.5, $S^{F'} = S^F$ and $F \cap S^F = F' \cap S^{F'}$ and $\alpha \mid\!\sim \gamma$ iff $\alpha \wedge \beta \mid\!\sim \gamma$. ∎

Karl Schlechta [11] has found a cumulative consequence relation $\mid\!\sim$ that cannot be defined by any stereotypical reasoning system yielded by a distance that satisfies Equation 5.2. The exact characterization of those cumulative relations that can be defined by stereotypical systems that satisfy Equation 5.2 is open. In the next section, we shall discuss another basic logical property of nonmonotonic system described in [2], Or, i.e., preferentiality.

6 Preferentiality

Suppose each one of two information sets, F and F' enable us to conclude that some formula α holds: all elements of $F \cap S^F$ and all elements of $F' \cap S^{F'}$ satisfy α. Does this imply that the union $F \cup F'$ enables us to conclude α, i.e., do all elements of $(F \cup F') \cap S^{F \cup F'}$ satisfy α? The discussion of [2] explains why this seems to be a natural property to expect. For example, assuming we would conclude that a bird that lives in the country flies, and that we would also conclude that a bird that lives in a city flies. Must we conclude that birds that live either in the country or in a city fly? Stereotypical reasoning does not always satisfy this property for the following reason. I guess that natural common-sense reasoning does not either, for the same reason. Suppose β and γ describe very different situations, whose best stereotypes are different. It may happen, nevertheless that the same property α will be shared both by models of $\beta \cap S^\beta$ and of $\gamma \cap S^\gamma$. But the best stereotype for $\beta \vee \gamma$ may be very general and some models of, say, $\beta \cap S^{\beta \vee \gamma}$ may not satisfy α. Intuitively, if the reasons for concluding α from β are very different from those for concluding α from γ, there is little hope we shall be able to conclude α from the disjunction $\beta \vee \gamma$.

There is one interesting case, though, the case $S^F = S^{F'}$, in which the desired conclusion follows if we strengthen one of the assumptions above. Since the function e does not depend on its third argument, by Assumption three, we shall write it as a function of two arguments. Let us assume:

- *Assumption four*: $e(A \cup A', B) = \min\{e(A, B), e(A', B)\}$.

Clearly, assumption one already implies $e(A \cup A', B) \leq \min\{e(A, B), e(A', B)\}$ and assumption four implies assumption one.

THEOREM 6.1
Let assumptions zero–four be satisfied. If $S^F = S^{F'}$, then $S^{F \cup F'} = S^F$.

PROOF. Notice that we do not claim that the nonmonotonic consequence relation defined is preferential. Assume $S^F = S^{F'}$. We must show that, for any stereotype S, we have $d(F \cup F', S^F) \leq d(F \cup F', S)$. Then,

$$d(F \cup F', S_F) = e((F \cap S_F) \cup (F' \cap S_{F'}), S_F - (F \cup F')) \leq$$

$$e(F \cap S_F, S_F - F) = d(F, S_F) \leq d(F, S).$$

Similarly $d(F \cup F', S_F) \leq d(F', S)$ and

$$d(F \cup F', S_F) \leq \min\{d(F, S), d(F', S)\} = d(F \cup F', S).$$

The last equality stems from assumption four. ∎

Consider our examples above. In Example 5.1, the function d is constant and therefore satisfies condition four. The consequence relation, being classical, is in fact preferential. In Example 5.2, the function d proposed does not satisfy assumption four, nevertheless the relation defined is preferential. In Example 5.3, the function d satisfies assumption four. The consequence relation defined, being classical, is in fact preferential. In Example 5.4, assumption four holds, and therefore the conclusions of Theorem 6.1, but the consequence relation defined is *not* preferential.

7 Conclusion

A formal description of stereotypical reasoning has been provided. Under reasonable assumptions about the way stereotypes are attached to information sets, this model yields a cumulative system. The assumptions one–three concerning the distance between information sets and stereotypes may perhaps be tested experimentally. Preferentiality has been discussed, and found unplausible in general, but a more limited natural property has been put in evidence. Again, preferentiality should be tested experimentally. The conditions proposed above that imply good logical behavior are sufficient but not necessary. Other conditions may be more natural and also sufficient. The structure of the set S of stereotypes, in particular, has been left completely arbitrary. A reasonable assumption may be that this set has a tree structure: i.e., that if S and T are any two stereotypes such that the intersection $S \cap T$ is not empty, then $S \subseteq T$ or $T \subseteq S$.

Acknowledgments

I would like to thank the members of the different audiences that reacted to the material above, while it was in gestation, and in particular to Yuri Gurevich, Pierre Livet, Drew McDermott, Karl Schlechta and Moshe Vardi. This work was partially supported by the Jean and Helene Alfassa fund for research in Artificial Intelligence and by grant 136/94-1 of the Israel Science Foundation on "New Perspectives on Nonmonotonic Reasoning".

References

[1] Dov M. Gabbay. Theoretical foundations for non-monotonic reasoning in expert systems. In Krzysztof R. Apt, editor, *Proc. of the NATO Advanced Study Institute on Logics and Models of Concurrent Systems*, pages 439–457, La Colle-sur-Loup, France, October 1985. Springer-Verlag.

[2] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1–2):167–207, July 1990.

[3] George Lakoff. *Women, Fire, and Dangerous Things*. The University of Chicago Press, 1987.

[4] David Makinson. General patterns in nonmonotonic reasoning. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, Nonmonotonic and Uncertain Reasoning, pages 35–110. Oxford University Press, 1994.

[5] Marvin Minsky. A framework for representing knowledge. In Patrick Henry Winston, editor, *The Psychology of Computer Vision*, Computer Science Series, chapter 6, pages 211–277. McGraw Hill, 1975.

[6] Marvin Minsky. A framework for representing knowledge. In John Haugeland, editor, *Mind Design*, chapter 3, pages 95–128. MIT Press, Cambridge, Mass., 1981.

[7] Hilary Putnam. *Mind, Language, and Reality.*, volume 2 of *Philosophical Papers*. Cambridge University Press, Cambridge, 1975.

[8] Eleanor Rosch. Natural categories. *Cognitive Psychology*, 4:328–350, 1973.

[9] Eleanor Rosch. Cognitive reference points. *Cognitive Psychology*, 7:532–47, 1975.

[10] Eleanor Rosch. Prototype classification and logical classification: The two systems. In E. Scholnick, editor, *New Trends in Cognitive Representation: Challenges to Piaget's Theory*, pages 73–86. Lawrence Erlbaum Associates, Hillsdale, N.J., 1983.

[11] Karl Schlechta. Remarks to stereotypical reasoning. personal communication, November 1997.

Received 24 October 1997. Revised 7 December 1997