

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Defaults Revisited Or "Tell Me If You'Re Guessing."

Permalink

<https://escholarship.org/uc/item/8qk591dt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 4(0)

Author

Nutter, Jane Terry

Publication Date

1982

Peer reviewed

DEFAULTS REVISITED
OR
"Tell me if you're guessing."

Jane Terry Nutter
Computer Science Department
SUNY at Buffalo
Buffalo, New York

ABSTRACT

This paper discusses default reasoning, distinguishing generalizations associated with defaults from both universals and statistical generalizations. I argue that conclusions based on defaults should be reported differently from conclusions which do not involve default reasoning, and that however we represent them, the related inference system must distinguish default claims from other propositions and treat them differently. Two existing analyses of default reasoning are briefly criticized in light of the distinctions presented.

1. Introduction.

A great deal of knowledge seems to take the form of generalizations: neither genuine universals, true of all things in their understood domains, nor simple statistical claims of "more than half", but claims which, although understood sometimes to fail, nevertheless warrant presumptions in the absence of conflicting information. Such generalizations are usually represented by defaults. This paper examines default generalizations, distinguishing them from universals and statistical claims, and pointing out some pitfalls their implementation presents. Especially, it behoves us to realize that answers based on default reasoning represent educated guesses; and however useful they may be, guesses cannot safely or honestly be handed out as facts.

2. Generalizations vs. universal and statistical claims.

In English, "all" rarely means "every single thing without exception", and failing to note this can produce unfortunate results [2]. To use Brachman's example, if we say that all elephants are four-legged gray mammals, and if we treat "all" as indicating genuine universality, then we have no way to talk about Clyde the unfortunate amputee elephant with only three legs. But suppose we always treat "all" as indicating a generalization Clyde the three-legged elephant, but unfortunately we can talk with equal ease about Clyde the non-mammalian elephant, or even about Clyde the non-elephant Indian elephant.

Generalizations cannot be treated like statistical claims either, although the difference here is more subtle. Most people realize that over half the population is female. Yet in the absence of information concerning a person's sex, one does not typically presume that the person in question is female (indeed, the presumption tends to go the other way!). By contrast, the number of flightless birds (emus, ostriches, kiwis, penguins, baby birds, etc.) is hardly negligible. Yet we feel justified in assuming of birds in general that they fly.

Generalizations usually represent causal claims, albeit masked and incomplete ones. Most birds fly, because the features which distinguish something as a bird evolved to facilitate flight. By contrast, statistical claims are evidence for, rather than

embody, causal claims. Furthermore, we accept statistical evidence as supporting causal claims only when there is independent reason to suppose that the phenomena involved are relevant to one another.

For example, I recall reading somewhere that for many years, the membership rolls of a baker's union in New York City precisely paralleled the births and deaths in a town in India. Whether this actually happened is not important here; my point is, it could well happen, and if it did, no reasonable person would take it as anything more than a striking (and somewhat humorous) coincidence.

The transitivity of inferences based on generalizations again distinguishes them from statistical claims. Presumability can be inherited through truth-functional inferences; but statistical relationships are far more complex, and statistical inferences follow utterly different rules.

For instance, consider the result of conjoining two statistical claims S and S' . Say the probability of S is x , and that of S' is y . Now what is the probability of $S \& S'$? Well, let's look at some examples.

Suppose the subject is coin tossing. Say S says "Toss 1 will be heads," and S' says "Toss 2 will be tails." Then $x = y = .5$, and the probability of "Toss 1 will be heads and toss 2 will be tails" we know to be .25, or xy . But is this always the case? Clearly not. Let S be as before, and let S' be "Toss 1 will be tails." Now the probability of $S \& S'$ is 0. If S is the same as S' , then the probability of the conjunction is the same as the probability of S .

Furthermore, statistical analyses tend to be applied to two fundamentally different sorts of situation. In the first kind, the various events are ex hypothesi independent of one another. We assume that the result of toss 1 does not affect the result of toss 2. In the second, a causal relationship is being sought or presumed. At this point, probabilities become inextricably linked to the theoretical context, and in some sense take on a different meaning. Given one set of results R , the probability of R will differ depending on the hypothesis relative to which it is computed. More importantly, what changes tends to be not the probabilities of individual occurrences, but precisely the probabilities of cooccurrences: that is, the probability of conjunctions changes, without that of the conjuncts changing. So no general rule captures the way the probability of a conjunction relates to the probability of the conjuncts.

3. Examples of default assumptions.

Suppose we are designing a "travel agent" system. The classical example of a default rule in this context is the assumption that, all else being equal, all trips originate wherever the customer currently is. This seems reasonable enough, but the system need hardly assume it, since it can request that information with no great loss of convenience.

But consider the following "rule": within the departure time limits the customer supplies, more direct connections are to be preferred over less direct ones. If someone says, "I'd like a round trip to New York," for the system then to ask, "Where are you leaving from?" seems reasonable; for it to ask, "Would you rather get there in one hour or nineteen and a half?" does not.

Furthermore, imagine a system which mindlessly produced every set of connections from Buffalo to New York — direct, via Albany, via Houston, via Seattle, via London, via Buenos Aires.... While the list might not go on forever, it will surely go on long enough to prove inconvenient. Some presumption must be made to order the alternatives so that reasonable ones get listed early.

Yet we cannot simply add a universal rule that direct routes are to be preferred over indirect ones, because it isn't always true. For example, some people refuse to use certain airlines or airports under any circumstances. Others will want to stop over for a few hours in some intermediate city.

There is also a more general problem. All else being equal, the cheaper of two routes is usually preferred over the more expensive. While the more direct route is usually also the cheaper, it is not always so. One can currently fly from Buffalo direct to Albany, which is shorter and more direct than changing flights in New York City. But it turns out that flying via New York is cheaper. Whether the customer wants to fly direct or via New York City will now depend on which is more important to the customer, convenient time scheduling or low price.

Hence the system cannot presume absolutely either that the more direct route is preferred, or that the cheaper route is. A guarded answer which presumes either, but with explicit reservations, will prove more useful than either a flat presumption which cannot be overruled (a universal) or a failure to make any presumption at all.

Other examples abound. If a customer asks to travel from New York to Cincinnati via Athens, we want the system to recognize that the customer probably means Athens, Ohio, and not Athens, Greece, or even Athens, Georgia. At the same time, this assumption should somehow be reflected in the system's response, lest travellers who mean to go to Athens, Georgia learn of Athens, Ohio by finding themselves there.

4. Problems defaults raise.

Perhaps the most common kind of default takes the form, "In the absence evidence that $\neg p$, you may infer p " [7,8]. When the system is asked " p ?" and finds the default rule, it attempts to derive $\neg p$. If it fails to do so, it returns p as the answer. Hence systems augmented by this kind of rule can take advantage of generalizations of the kind above. So far, so good.

But this procedure only looks reasonable so long as we deal with questions like "Can Roger the bird fly?" Then, saying "Of course, he's a bird," seems unobjectionable — but only because nothing depends on the answer. Notice that if we don't care what the answers to our questions are, there is little reason to implement defaults. After all, if we don't care, we can as well say "I don't know" as either yes or no.

But suppose that we do care what answer we get. For instance, consider a medical diagnostic and treatment-recommending system. Suppose that for a particular set of symptoms, treatment x is generally very beneficial, but that in the exceptional cases treatment x invariably kills. Now if A has the symptoms in question, surely we do

not want to recommend treatment x solely on the grounds that we don't yet know that A is exceptional. On the other hand, if the symptoms in question can themselves prove fatal, nor do we want to say we don't know anything about what to do for A .

In this kind of case, we would like the system to say something like, "Treatment x usually helps," or "Presumably treatment x helps." Even better would be an answer which directly tells the user what the counterindications are; but at the very least, a responsible system should warn the user that the information results from a presumption, and not an inference. Once the system has issued the warning, the user can then pursue it in further questions.

A further difficulty with defaults lies in deciding what it means for them to be true or false. Clearly "If Roger is a bird, then presumably Roger can fly" can be true even if Roger is a bird, but Roger can not fly. Indeed, "Presumably Roger can fly" can be true, even though "Roger can fly" is false. That is the whole point of saying "presumably": it protects the speaker from saying something false when the facts go the "wrong" way. That is what it means to give a guarded response.

Hence the truth value of defaults cannot be a simple function of the truth values of their component propositions: default operators are not truth functional. Furthermore, defaults make sense because they reflect causal (and hence non-logical) connections among their constituents. The missing information guarantees that their content cannot be a simple function of the contents of the components. But then we should not expect to be able to give a purely logical account of defaults [4].

5. Problems with two proposed solutions.

Several approaches to defaults have been suggested. Some researchers treat defaults as modalized [6,7,8]. Several problems with this approach have been pointed out already (see e.g. [3]). In addition, this approach interprets "In general, birds fly" as something like "If x is a bird and it is compatible with what we know that x flies, then x flies" [7]. But this is only true if every single bird without exception which we do not know to be flightless does in fact fly. That is, if McDermott's version of the generalization is true, it can never be the case that some bird does not fly and we can not prove that it doesn't. But this is surely not what the generalization means.

The fuzzy logic approach [1,5,9,10] uses a continuum of truth values in the closed range $\langle 0,1 \rangle$ instead of simply "true" and "false". Several questions immediately arise. First, every "assertion" in the data base must have an associated truth value; where are we to get these from? Second, how are the truth values of propositions related to those of their components, and how are the truth values of conclusions related to those of the premises of the demonstration in question? Preliminary results [1] boil down to the unsurprising claim that the conclusions are no better than the premises, but also on the whole no worse (where "better" is interpreted as numerical "greater than"). It is significant that this is already non-trivial to establish. Third, how do we deal with the apparent result that different demonstrations of the same proposition "establish" different truth values?

But the largest problem, in my opinion, lies in the irresistible temptation to view these fuzzy truth values as probabilities. This tendency is encouraged by the need to assign what, in context, look much like Bayesian prior probabilities to the

propositions in the data base. Some kind of Bayesian analysis may prove useful in A.I. systems; but there is no "cut-rate" way of doing it. Neither fuzzy logic nor default reasoning adequately analyzes probability. Under the circumstances, it seems best to avoid a system which misleads to this extent.

6. Conclusion.

We would like some way to deal with the "funny" truth status of default rules and of conclusions drawn on the basis of default assumptions; but neither modality nor fuzzy truth values seems to capture the desired effect. Furthermore, there seems good reason to suppose that no purely logical analysis could.

But this does not rule out the possibility that logical restrictions on defaults and their consequences can be found and described, on the basis of which a system of inferences allowing default reasoning can be developed. We are currently developing a semantics for default reasoning which treats defaults as propositional operators and which we hope will provide such a basis. Once this has been done, we can hope to deal with defaults in a reasonable and useful way.

Hence an A.I. system which deals with defaults successfully must also have at least two properties which existing proposals lack. First, it must delineate the logical restrictions on defaults and their consequences without ruling out the existence of genuine exceptions, i.e., recognizing that default reasoning sometimes gives the wrong answer. In doing so, it should be careful to distinguish default generalizations both from genuine universals and from statistical generalizations. And second, when the system gives answers which are based on default reasoning, it should admit this weakness by issuing warnings with them. For without such warnings, default reasoning by any scheme is not only unsound: it is also unsafe.

8. Acknowledgments.

I would like to thank Stuart Shapiro and the members of the SNePS Research Group at SUNY/Buffalo for their many helpful comments and suggestions.

8. References.

- [1] Aronson, A.R., Jacobs, B.E., and Minker, J. A note on fuzzy deduction. JACM v. 27 (1980) 599-603.
- [2] Brachman, R.J. "I lied about the trees" or defaults and definitions in knowledge representation. Draft (1982).
- [3] Davis, M. The mathematics of non-monotonic reasoning. A.I. v. 13 (1980) 73-80.
- [4] Israel, D.J. What's wrong with non-monotonic logic? Proc. First Annual National Conference on Artificial Intelligence, American Association for Artificial Intelligence (1980) 99-101.
- [5] Lee, R.C.T. Fuzzy logic and the resolution principle. JACM v. 19 (1972) 109-119.
- [6] McDermott, D.V. and Doyle, J. Non-monotonic logic I. A.I. v. 13 (1980) 41-72.
- [7] McDermott, D. Non-monotonic logic II. JACM v. 29 (1982) 33-57.
- [8] Reiter, R. A logic for default reasoning. A.I. v. 13 (1980) 81-132.
- [9] Zadeh, L.A. Fuzzy sets. Inf. Control v. 8 (1965) 338-353.
- [10] Zadeh, L.A. Fuzzy algorithms. Inf. Control v. 12 (1968) 92-102.