CS-TR-3280 | May 26, 1994

UMIACS-TR-94-64

What Experts Deny, Novices Must Understand

Michael Miller

Don Perlis*

*Institute for Advanced Computer Studies

Computer Science Department

AV Williams Bldg

University of Maryland

College Park, MD 20742

mm@cs.umd.edu

perlis@cs.umd.edu

Abstract

We consider the problem of representing the denial of default information. We show that such denials are important parts of commonsense reasoning. Moreover, their representation is not a simple matter of negating traditional representations of default information. We have found a solution by separating default information into use and trend portions. This approach may also afford a more compact way to represent defaults in general.

Work done with partial support from NSF grant IRI9311988.

What Experts Deny, Novices Must Understand*

Michael Miller and Donald Perlis

Department of Computer Science

University of Maryland

College Park, MD 20742

(301) 405-{2716,2685}

{mm,perlis}@cs.umd.edu

*This research was supported in part by NSF grant IRI9311988

1 Introduction

The formal study of inference has undergone an explosion over the past 15 years with the introduction of nonmonotonic (default) reasoning. Here we consider some ramifications of default reasoning that may be applicable to expertise. It is now well recognized that a great deal of practical knowledge about any domain is in the form of defaults. This suggests the possibility of characterizing expertise in terms of defaults and default use. We will do this toward the end of the paper after developing some prerequisite machinery. One ingredient in our characterization is the capacity to deny defaults; this has received little attention in the literature and seems to require some new technical devices.

1.1 Default reasoning

The commonsense world is far too complex for reasoners, human or otherwise, to be aware of all information that may be salient to a situation at any given time. Therefore very often a conclusion one draws is not necessarily true; additional information might have "defeated" the conclusion and even prevented the inference in the first place. When I learn that Tweety is a bird, it is reasonable for me to conclude that Tweety can fly, since birds *typically* fly; but the conclusion may be false nevertheless. An endless supply of counterexamples can be constructed: Tweety may have had her wings clipped or may have become a victim of an oil-spill or may be a penguin; all good reasons for her not to be a flyer. Moreover, had I known the additional information that Tweety was a penguin, I would not have concluded that she can fly in the first place. Very many cases of practical reasoning are of this sort. Yet this form of reasoning lies outside the framework of traditional "monotonic" logic in which more information leads to more (not fewer) conclusions [10].

Researchers have addressed this issue by developing formalisms for default (or nonmonotonic) reasoning, in which something inferred from one set of information may fail to be inferred from that set augmented with additional information. Three well-known such formalisms are Reiter's default logic (DL) [12], McDermott and Doyle's non-monotonic logic (NML) [8], and McCarthy's circumscription [7]. Each of these offers a formal treatment of a reasoner coming to a reasonable yet defeasible conclusion, based on whatever knowledge is available plus some default rule(s). Just what *is* a reasonable conclusion, how it is reached, and what the default rules are, vary from formalism to formalism.

In looking more closely at the Tweety example, we can see two elements at work. The first is a default principle (which we will usually refer to simply as a *default*) about typical features of birds, such as "birds typically fly", and the second is a means to employ that knowledge to draw a conclusion about particular birds (e.g., Tweety). Most of the research in default reasoning has dealt with both of these issues in varying ways. It is our contention in this paper that for some purposes, such as representing novice-expert differences, it is important to sharply separate these two aspects. As a matter of notation we will write $Bird(x) \stackrel{typ}{\rightarrow} Fly(x)$ for the default "birds typically fly"; more generally we write $\Psi \stackrel{typ}{\rightarrow} \Phi$ for "Ψ's typically are Φ's". This notation is intended to be neutral among all the standard default formalisms: in DL, $\stackrel{typ}{\rightarrow}$ would be expressed as an inference rule, in NML it would involve a consistency predicate, and in circumscription it would involve an abnormality predicate.

1.2 Expertise

Cognitive scientists have intensively studied expertise, in many settings. Not surprisingly it is novice-expert differences that motivate many of these studies. Mayer [6] describes several knowledge-based distinctions between experts and novices including: (i) novices tend to store their knowledge in small fragmented units, while experts store theirs in larger interconnected functional units and (ii) experts tend to know and use deep *structural* knowledge while novices use more *superficial* knowledge.

Experts have substantial knowledge about their domain of expertise, yet their judgements are based on surprisingly little information (see [2] and [3]).[1] For instance in medical diagnoses, experts form a few, well-focused hypotheses early in the process (even before adequate data is presented), and use that to guide what data they seek [5]. In contrast novices

[1]While one might suppose that the expert knows which relevant pieces of information to pick out and which pieces to ignore, there is some evidence indicating that this is not the case. To the contrary, it has been shown that experts are often influenced by irrelevant information and as a result their decisions may turn out to be incorrect or unreliable [13]. Moreover, experts are often unable to provide convincing accounts of how they make their judgments (see [1]).

follow more of a broad searching strategy, don't have well-formulated hypotheses early on, and obtain much more information (much of which is not of value).

Feltovich et al [4] suggest several postulates regarding different characteristics of novice and expert knowledge bases (in the domain of medicine). Included is that the novice's knowledge base is sparse—it lacks the cross referencing and clustering structure of the expert's dense knowledge base. Also, the novice's knowledge is more imprecise than the expert's. This is, in part, due to the expert's "fine-tuning" of her knowledge through clinical experience. These differences result in novice expectations (in the field of medical diagnosis) being "either overly general, allowing clinical findings that should not occur, or overly specific not allowing the legitimate range". Rau [11] provides a computational treatment of breadth and salience of knowledge which appears related to the appropriate range of application of knowledge in making judgments.

Both experts and novices draw conclusions based on their knowledge. They may differ in their knowledge; e.g., an ornithologist will likely believe the default that cardinals are either red (males) or russet (females)—with rare exceptions being white (albinos)—while a casual observer (i.e., a novice) may mistakenly believe simply that cardinals are typically red, a narrower range of possibilities than that used by the expert.[2] They may also differ in how they draw conclusions from their knowledge. This knowledge/conclusion-drawing distinction can be misleading, though. How one draws conclusions can depend on one's knowledge, i.e., the inference procedures can be partly encoded as declarative knowledge, such as in typicality statements; this will become important as our discussion proceeds.

1.3 An example

Imagine Tommy to be a (novice) robot and Sue an (expert) ornithologist responsible for training Tommy. Sue takes Tommy to the zoo and the following conversation ensues:

Tommy: "Look, someone's feeding all those birds!"
Sue: "Do you know what kind of birds they are?"
Tommy: "No."
Sue: "They're cardinals."

[2]We conjecture that such "default ranges"—which will be addressed in what follows—are important aspects of the phenomenon studied in [4] and [11].

Tommy: "But they're not red. I thought cardinals were supposed to be bright red."
Sue: "I used to think that also. Quite a few are red but many others are russet like these."
Tommy: "Are all cardinals red or russet?"
Sue: "Almost; a very few are white albinos. I wonder why only russet ones are here."
Sue (to the zookeeper): "Why are there only female cardinals here?"[3]
Zookeeper: "We put the males in another cage for an experiment."

From the dialog we can presume that in the past Tommy has picked up a bad default: that cardinals typically are red. Sue, the expert, teaches him otherwise. This requires that her language—as well as Tommy's—allow for the expression of the *denial* of defaults (e.g., roughly, "it is not the case that cardinals typically are red"). There are a number of other features illustrated in the above dialog which we will point out in sections 3 and 4. First, in Section 2 we deal with the denial of defaults.

2 Range defaults and default denials

2.1 Range defaults

In the above dialog there is evidence that Sue holds a special kind of default called a *range default* about cardinal color. Specifically, Sue believes that cardinals typically are red *or* russet, though there are exceptions—the albinos—and she denies that cardinals typically are red. Moreover she would also deny that cardinals typically are russet: there are proportionately too many red and too many russet cardinals for these latter defaults to be sensible. Here's another example: People typically are male or female[4]—yet there are too many male *and* too many female people to exclude either maleness or femaleness as typical. Both of the defaults "people typically are male" and "people typically are female" are too restrictive and hence inappropriate.

More formally, a range default is an accepted default of the form "P's are typically Q's" where Q is a disjunction (i.e., $\bigvee_{i \in I} \Phi_i$) and for every shorter disjunction S formed from the (disjunctive) components of Q, the (sub-) default "P's are typically S's" is denied. Thus a range default has two parts: an affirmed disjunctive default (e.g., "cardinals typically are red *or* russet") and various

[3]Sue, but not Tommy, knows that female cardinals are russet and males are red.
[4]Note that this too is a default, not a universal fact. There are rare hermaphroditic or neuter persons who are neither (singularly) male nor (singularly) female; and perhaps gender-changes can be similarly construed.

denied defaults (e.g., "it is not the case that cardinals typically are red" and "it is not the case that cardinals typically are russet"). (This will be further formalized in Section 2.2.) The above examples above suggest that some of the detailed, fine-tuned, substantial quantity of expert knowledge takes the form of range defaults as opposed to simple defaults.[5]

Why bother with range defaults, or more precisely, with default denials at all? Knowing that a particular bird is a cardinal and that typically cardinals are red or russet, the latter being the affirmative disjunctive part of a range default, should lead one to conclude that the bird is red or russet, unless she knows to the contrary. A reasoner need not go on to deny the individual defaults "typically cardinals are red" and "typically cardinals are russet" in reaching that conclusion. Why do range defaults deserve special mention? There are two reasons: (i) their representation presents a challenge (see [9]) and (ii) they can play a role in reasoning that is not accomplished solely with "ordinary" defaults. For one there are cases in commonsense reasoning where it is not only important to reach the correct default conclusion, but also to have meta-knowledge about one's own defaults which itself can be reasoned with and about.

For instance in the above dialog Sue notices that the collection of birds in the zoo is unusual: they are all russet and hence female. But the plain disjunctive default, namely that cardinals typically are red or russet, does not prompt this conclusion. She does have excellent reason to think it is an oddity, though, because she has the additional information that it is *not* the case that cardinals typically are russet. She uses this to form the observation that an *unusual* collection of cardinals has gathered at the zoo, and hence to wonder: Why have only female cardinals gathered? A novice, on the other hand, without Sue's knowledge, would have no reason to inquire further.

This sort of knowledge that an expert may have regarding cardinals is precisely what a range default about cardinal color expresses, and this knowledge is crucial to the reasoning illustrated. Thus, not only is the formal representation of range defaults of interest in a purely theoretical sense (Can a formalism represent them?), it also has pragmatic ramifications for commonsense reasoning formalisms.

The representation and use of range defaults clearly hinges on the representation and use of denied defaults: defaults that not only are not part of one's belief base but are explicitly believed to be false. Denying a default requires explicit recognition of that default as a mistake to avoid.

[5] In a sense we all become less "novice" and more "expert" about everyday matters concerning people, relationships, gender, etc. as we grow up.

This we suggest is abundant in an expert's repertoire and less so in a novice's.[6]

2.2 The problem of denying defaults

How can we formally represent range defaults? From our previous discussion we need two things: firstly a plain disjunctive default principle

$$\Psi \overset{typ}{\rightarrow} \bigvee_{i \in I} \Phi_i \tag{1}$$

and secondly assertions to the effect that the range cannot be restricted any further. That is, for every non-empty proper subset J of I the associated potential default principle is *denied*:

$$\neg(\Psi \overset{typ}{\rightarrow} \bigvee_{i \in J} \Phi_i) \tag{2}$$

Thus (1) and (2) together formalize the total range default.

Notice that (1) is doing double-duty.[7] On the one hand it encodes a typicality statement about the population of Ψ-things (i.e., a statement asserting a "population trend") and on the other it encodes a means of drawing a conclusion about a specific Ψ-thing. But then (2) is puzzling: Which is being denied, the typicality statement or the inference procedure? If we take denial literally and use classical negation—as indicated in (2)—then a problem surfaces immediately in DL where defaults are written as inference rules, since there is no recognized formal notion of the negation of a rule of inference. In NML and circumscription, negating a default results in a *counterexample* axiom, which simply records that the default has led to an error in at least one case and is *not* an assertion that the default itself is a bad one that should not be used. (We analyze these issues in detail in [9].)

In short, default denial seems to create varied difficulties for standard formalisms. We will utilize the approach of [9] which appears to handle all of the problems simultaneously and uniformly across all the formalisms. We summarize this briefly in Section 2.3 below and then apply it our sample expert-novice dialog, in Section 3.

[6] We are unaware of any empirical studies bearing on this, however.

[7] The reader may wish to view $\overset{typ}{\rightarrow}$ in terms of a particular formalism, say DL, to see the dual functions served by (1). We will illustrate this below in Section 2.3.

2.3 Formally separating the two default features

Defaults tend to be formally represented in combined form containing a conflation between a general typicality (or trend) about a population on the one hand, and a sanctioning of inferred conclusions about a particular member of the population on the other hand. The negation of such a representation (when possible at all) then mixes together both the unsoundness of the inferences, and the denial of the population-trend, even if (as in the case of range defaults) only the latter is wanted.

The key, then, is to view a default as having two complementary features. One is the typicality statement Typ giving default information as a trend about the commonsense world. The other is the inferential mechanism Inf by means of which Typ is used to produce default conclusions. What is needed for default denial is negating Typ, not Inf. But when Inf and Typ are combined in a single representation, this is problematic.

Our approach is as follows: we treat standard default mechanisms as largely playing the role of Inf, and adjoin a separate typicality statement to play the role of Typ. Should we want to deny the trend itself we assert $\neg Typ$, the negation of the latter statement.

As an example, in DL we first replace each default rule of the form[8]

$$\frac{\Psi \ : \ \Phi}{\Phi} \tag{3}$$

by the axiom

$$Typically(\Psi, \Phi) \tag{4}$$

and adjoin a single new inference rule

$$\frac{P \wedge Typically(P, Q) \ : \ Q}{Q} \tag{5}$$

where P and Q are second-order variables that can be bound to first-order expressions such as Ψ and Φ. Then (4) and (5) together produce the effect of (3) by binding P to Ψ and Q to Ψ. For an individual a, if (it is believed that) $\Psi(a)$, then $\Phi(a)$ is concluded when consistent. On the other hand, for a particular Ψ_0 and Φ_0, if $\Psi_0(a)$ and $Typically(\Psi_0, \Phi_0)$ is *not* believed then (5) does not produce $\Phi_0(a)$. Therefore it is the presence or absence of statements like (4) which control the use

[8] We remind the reader that $\frac{\Psi \ : \ \Phi}{\Psi}$ has the approximate reading: If Ψ is believed, and if Φ is consistent with all that is believed, then Φ is inferred (believed).

of (5). For a particular Ψ_0 and Φ_0, we can even explicitly deny such a population trend by asserting

$$\neg Typically(\Psi_0, \Phi_0) \tag{6}$$

To return to the issue of expertise, the above discussion lends itself to the idea that a single general-purpose default-use mechanism, as in (5), may suffice for experts and novices alike.[9] We suggest that default reasoning differences between experts and novices may hinge on differences in knowledge of the form (4) as well as denials (6); we pursue these ideas in Sections 3 and 4.

3 Example revisited—pieces of a formal treatment

Once typicality statements are separated from the inferential mechanism range defaults can be accurately represented. We illustrate this by casting fragments of the knowledge evident in our sample dialog, into a more formal setting.

At the start of the conversation, Tommy's and Sue's knowledge bases (KB) include:

Fragment of Tommy's KB

$Typically(Cardinal, Red)$
$(\forall x)\ (Cardinal(x) \rightarrow Bird(x)$
$Typically(Bird, Fly)$

Fragment of Sue's KB

$Typically(Cardinal, Red \vee Russet)$
$\neg Typically(Cardinal, Red)$
$\neg Typically(Cardinal, Russet)$
$(\forall x)\ Cardinal(x) \rightarrow Bird(x)$
$Typically(Bird, Fly)$
$Typically(Cardinal \wedge Female, Russet)$
$Typically(Cardinal \wedge Male, Red)$
$(\forall x)\ Cardinal(x) \rightarrow Red(x) \vee Russet(x) \vee White(x)$

Notice that both Sue and Tommy hold a default about cardinal color, but Sue's is a range default (contained in the first three beliefs, above) whereas Tommy's is not. Additionally Sue holds simple defaults concerning female cardinal color and male cardinal color and a firm (non-default) belief about all possible cardinal colors; Tommy does not.

[9]We do not suggest that *all* reasoning mechanisms are necessarily shared by experts and novices, rather we speculate that a specialized default mechanism is.

Using the above knowledge bases, the following conclusions are readily drawn (in DL, say) about an individual cardinal b_0, given the additional fact $Cardinal(b_0)$. (Starred $(*)$ conclusions are default inferences.)

Tommy's Initial Conclusions	**Sue's Initial Conclusions**
$*Red(b_0)$	$*Red(b_0) \vee Russet(b_0)$
$Bird(b_0)$	$Bird(b_0)$
$*Fly(b_0)$	$*Fly(b_0)$
	$Red(b_0) \vee Russet(b_0) \vee White(b_0)$

From the still further information $Male(b_0)$, Sue would conclude $Red(b_0)$; and from $Female(b_0)$ she would conclude $Russet(b_0)$. Tommy would conclude nothing further if given either $Male(b_0)$ or $Female(b_0)$.

At the end of the dialog, after Sue's instruction, Tommy's knowledge base still contains a single default about cardinal color, but now it is a range default:

Fragment of Tommy's New KB

$Typically(Cardinal, Red \vee Russet)$
$\neg Typically(Cardinal, Red)$
$\neg Typically(Cardinal, Russet)$
$(\forall x)\ Cardinal(x) \rightarrow Bird(x)$
$Typically(Bird, Fly)$

Tommy still lacks some of Sue's knowledge about cardinal color, namely $Typically(Cardinal \wedge Female, Russet)$ and $Typically(Cardinal \wedge Male, Red)$, but his previously held (bad) default is now gone. From $Cardinal(b_0)$ Tommy now can conclude $Red(b_0) \vee Russet(b_0)$.

4 A tentative formal characterization of expertise

We can use our foregoing analysis and sample dialog to motivate a characterization of expertise in terms of a default-reasoning framework. Recall:

(1) **Tommy:** "Look, someone's feeding all those birds!"
(2) **Sue:** "Do you know what kind of birds they are?"
(3) **Tommy:** "No."
(4) **Sue:** "They're cardinals."
(5) **Tommy:** "But they're not red. I thought cardinals were supposed to be bright red."
(6) **Sue:** "I used to think that also. Quite a few are red but many others are russet like these."
(7) **Tommy:** "Are all cardinals red or russet?"
(8) **Sue:** "Almost; a very few are white albinos. I wonder why only russet ones are here."
(9) **Sue (to the zookeeper):** "Why are there only female cardinals here?"
(10) **Zookeeper:** "We put the males in another cage for an experiment."

Sue (the expert) enters the scenario with more defaults than Tommy (the novice), including a range default about cardinal color. Her defaults are presumably well-grounded in breadth, i.e., many observations (not only her own, but also of others who have taught her) have led her to believe them. She also has default denials; knows about exceptions to her (range) defaults, e.g., albino cardinals; knows that the collection of female cardinals is unusual or abnormal; and knows when she does *not* know enough to explain the unusual collection of cardinals.

We propose the following tentative formal characterization of expertise:

- **Experts hold a substantial number of defaults (compared to a novice) in the domain of expertise.** Sue holds more, but related, defaults than does Tommy, even after she has given him some instruction, at which time Tommy still remains unaware that female cardinals typically are russet, and males typically are red.
- **Expert's defaults are well grounded in breadth.** (See [11] for a discussion of breadth of knowledge and [4] for related issues.)
- **Novices' defaults tend to be overly specific or overly general.** Tommy's original

default concerning cardinal color, that "cardinals typically are red" (indicated by line 5 of the dialog), is too narrow; Sue provides a more accurate range (red or russet) for his default (line 6).

- **Experts also hold a substantial supply of default denials, each indicating a common mistake to be avoided.** Sue's statement "Quite a few are red but many others are russet" (line 6) indicates that she also holds the default denial that "it is not the case that cardinals are typically red" as well as the default denial that "it is not the case that cardinals are typically russet".
- **Experts have the ability to deal with individual exceptions to defaults.** Sue knows that some exceptional cardinals are white (line 8).
- **Experts know when collections of things are abnormal, based on combined defaults and default denials.** Sue wonders why only russet cardinals have gathered at the zoo (line 8)—an abnormal situation given her default ("cardinals are typically red or russet") together with her default denial ("it is not the case that cardinals are typically russet".)
- **Experts know when they don't know, i.e., they know when it's time to ask for help.** As Sue does in line 9.

5 Conclusion

We have examined expertise from the point of view of default reasoning, and suggested a number of connections. In particular, we noted the importance of the ability to *deny* a default principle, and illustrated a formal mechanism for this in the context of a dialog between an expert and a novice. While we do not claim that defaults (and denials) are all there is to expertise, we do think that this is a fruitful dimension along which to explore expert reasoning, including novice-expert shifts.

We did not give a formal treatment of the entire Tommy-Sue dialog. That would have involved many additional mechanisms beyond the scope of this paper, and indeed beyond the state of the art in automated reasoning. Among other things, appropriate means would be needed for reasoning about groups, including statistical and set-theoretic aspects, in ways compatible with default reasoning; also time, change-of-mind, advice-taking and learning would enter importantly. All of these and more would have to be combined into a single robust system to achieve a satisfactory automated reasoner like Tommy (or Sue). While there is much progress on these various themes, integrating them has not to our knowledge even been attempted. Such an integration is a major research goal of our on-going work.

References

[1] W. J. Clancey. The situated cognition perspective on knowledge and context, 1993. Presented at the Third International Workshop on Human and Machine Cognition, Seaside, FL.

[2] E. Ebbesen and V. Konechi. Decision making and information integration in the courts: The setting of bail. *Journal of Personality and Social Psychology*, 32:805–821, 1975.

[3] H. Einhorn. Expert judgment: Some necessary conditions and an example. *Journal of Applied Psychology*, 59:562–571, 1974.

[4] P. Feltovich, P. Johnson, J. Moller, and D. Swanson. LCS: The role and development of medical knowledge in diagnostic expertise. In W. Clancy and E. Shortliffe, editors, *Readings in Medical Artificial Intelligence: The first decade*, pages 275–319. Addison-Wesley, 1984.

[5] J. Kassirer and G. A. Gorry. Clinical problem solving: A behavioral analysis. *Annals of Internal Medicine*, 89:245–255, 1978.

[6] R. E. Mayer. *Thinking, Problem Solving, and Cognition.* W. H. Freeman and Co., NY, 2nd edition, 1992.

[7] J. McCarthy. Circumscription: A form of non-monotonic reasoning. *Artificial Intelligence*, 13(1,2):27–39, 1980.

[8] D. McDermott and J. Doyle. Non-monotonic logic I. *Artificial Intelligence*, 13(1,2):41–72, 1980.

[9] M. Miller and D. Perlis. Defaults denied. *In preparation*, 1994.

[10] M. Minsky. A framework for representing knowledge. In P. Winston, editor, *The Psychology of Computer Vision.* McGraw-Hill, 1975.

[11] L. F. Rau. A computational approach to meta-knowledge: Calculating breadth and salience. Technical report, 1993. GE AI Lab, Report 93CRD094.

[12] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1,2):81–132, 1980.

[13] J. Shanteau. Some unasked questions about the psychology of expert decision makers. In *Proceedings of the 1984 IEEE Conference on Systems, Man, and Cybernetics*, pages 408–412, NY, 1984.