

VARIABLE PRECISION LOGIC

by

Ryszard S. Michalski

Patrick H. Winston

Artificial Intelligence Memo., No 857, MIT, Cambridge, MA, September, 1985.
Accepted to AI Journal, 1986.

Artificial Intelligence Laboratory
Massachusetts Institute of Technology

MIT AI Memo 857

August 1985

VARIABLE PRECISION LOGIC

by

Ryszard S. Michalski and Patrick H. Winston

Abstract. *Variable precision logic* is concerned with problems of reasoning with incomplete information and resource constraints. It offers mechanisms for handling trade-offs between the precision of inferences and the computational efficiency of deriving them. Two aspects of precision are the *specificity* of conclusions and the *certainty* of belief in them; we address primarily certainty and employ *censored production rules* as an underlying representational and computational mechanism. These censored production rules are created by augmenting ordinary production rules with an *exception* condition and are written in the form *if A then B unless C*, where *C* is the exception condition.

From a *control* viewpoint, censored production rules are intended for situations in which the implication $A \Rightarrow B$ holds frequently and the assertion *C* holds rarely. Systems using censored production rules are free to ignore the exception conditions when resources are tight. Given more time, the exception conditions are examined, lending credibility to high-speed answers or changing them. Such logical systems therefore exhibit variable certainty of conclusions, reflecting variable investment of computational resources in conducting reasoning. From a logical viewpoint, the *unless* operator between *B* and *C* acts as the exclusive-or operator. From an expository viewpoint, the *if A then B* part of the censored production rule expresses important information (e.g., a causal relationship), while the *unless C* part acts only as a switch that changes the polarity of *B* to $\neg B$ when *C* holds.

Expositive properties are captured quantitatively by augmenting censored rules with two parameters that indicate the certainty of the implication *if A then B*. Parameter δ is the certainty when the truth value of *C* is unknown, and γ is the certainty when *C* is known to be false.

Acknowledgment. This research was done in part at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology and in part at the Department of Computer Science at the University of Illinois. Support for MIT's artificial-intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505. Professor Michalski's research at the University of Illinois is supported in part by the National Science Foundation under grant NSF DCR 84-06801, in part by the Office of Naval Research under grant N00014-82-K-0186, and in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research Contract N00014-K-85-0878.

Variable Precision Logic Is Concerned with Specificity and Certainty

You cannot tell an ordinary logic-based reasoning system much about how you want it to do its job. You cannot give the following instructions, for example:

- Give me a reasonable answer immediately, even if somewhat general; if there is enough time, give me a more specific answer.
- Give me a reasonable answer immediately; if there is enough time, tell me you are more confident in the answer or change your mind and give me another, better answer.

Suppose, for example, that you want to know what John is doing, given that it is Sunday. A quick answer is that he is *probably* working in his yard. A more *specific* answer, obtained by taking into consideration the time of year, is that he is *specifically* raking leaves. A more *certain* answer, obtained by noting nice weather, is that he is *certainly* working in his yard, rather than reading.

A system that gives more specific answers, given more time, is what we call a *variable-specificity system*. A system that gives more certain answers, given more time, is what we call a *variable-certainty system*. There can be various combinations of the two systems, reflecting the fact that specificity and certainty are inversely related: we can gain specificity at the expense of certainty, or we can gain certainty by sacrificing specificity.¹ This point can be illustrated by going back to our example about what John is doing on Sunday. In this example, the less specific statement, "John is working outdoors," is more certain than the more specific statement, "John is working in his yard."

Variable specificity and variable certainty are two aspects of what we call *variable precision*. Thus, in general, we say that a variable-precision system is a system that exhibits either variable specificity or variable certainty or some trade off between the two.

The purpose of this paper is to show how we have extended ordinary logic so as to enable logic-based systems to exhibit variable precision in which certainty varies while specificity stays constant. As a vehicle to implement such a logic system we employ *censored production rules*, which are production rules with exceptions.

¹ We use the terms *specificity* and *certainty* in ways that are reminiscent of the use of the terms *precision* and *accuracy* in measurement theory. However, *precision* has to do with the number of significant digits associated with a quantity, while *specificity* is the opposite of *generality*. For example, we say that the concept of a triangle is a more *specific* concept than the concept of a polygon, rather than a more *precise* one. Similarly, the term *accuracy* has to do with the numerical difference between the measured value and true value of a quantity, while *certainty* is a measure of confidence that a given statement represents a true statement.

Before we explain censored production rules in detail, however, we discuss the need for rule-repair mechanisms, which is our second motive for extending logic with censors.

Censors Make Exceptions Explicit, Facilitating Rule Repair

The rules that we postulate for the world are normally tentative because there are few regularities that hold universally for any entity or class of entities. There is always a possibility that rules may have to be revised in the face of new facts or new happenings. Heraclitus's *panta rei*, i.e., all is in motion, applies not only to the outside reality but also to our mental representations of that reality. Because our knowledge is fluid and subject to modifications, a representation of that knowledge should make modifications easy and natural. A corollary is that changes to formal descriptions should closely reflect changes in our own thinking.

What then do we do when a rule (or a theory) which has worked well in the past does not work in some newly observed situation? There are several possibilities:

- Action 1. To consider the rule invalid, and ignore it in the future.
- Action 2. To continue to use the rule without change, realizing that using it will result in error occasionally.
- Action 3. To modify the rule, so that the rule applies correctly to all encountered situations.
- Action 4. To develop a new rule, substituting the new rule for the old one.
- Action 5. To remember the situations for which the rule does not work, treating them as exceptions.

All of these choices force a trade-off between estimated cost and estimated benefit. In some situations we may not be able to afford the time or other resources to make modifications to the rule before we need to use it. Also, which type of rule repair is best depends on the type of contradiction found to the rule.

Action 1, to invalidate the rule, is simple and prevents us from making errors, but may leave us handicapped. If the rule worked well in many cases, then invalidating it deprives us of the benefit of using it when it does work.

Action 2, to use the rule without change, is also simple. It preserves the benefit of employing it when it does work, but using it will lead to some errors.

Action 3 calls for creating a new rule by modifying the old one, and action 4 calls for developing a new rule from scratch. If the modification to be made to a rule is small, then action 3 is the better choice. But if this modification is complicated or unclear, then action 4 is the better

choice. In general, both actions lead to a better and more precise rule, but both require time and effort. In science, where standards for precision and certainty are high, one of these two actions is the usual choice, independent of the cost. (The problem of incrementally refining rules to accommodate new facts is explored in Reinke and Michalski [1985].)

If the exceptions are few, then action 5, to remember exception conditions, is a good choice. It preserves the usefulness of the old rule, but prevents us from making mistakes in situations recognized as exceptions. Even when there are more than a few exceptions, remembering them still may be the best action to take, particularly when it is not clear how to make changes to the old rule or how to create a new one.

Another situation for which exception remembering may be the best choice is when a modified or completely new rule is significantly more complicated than the original rule. A simple rule with exceptions may be better than a complicated one without exceptions, particularly when the exceptions occur only rarely. When the number of exceptions grows, they may be generalized via the introduction of a new rule, reducing the overall complexity of the original rule.

From this point of view, the purpose of this paper is to introduce ideas centered on the exception-remembering approach to knowledge modification. Again, we employ *censored production rules*, which are production rules with exceptions. This leads us to forms of representation that we believe are more natural and comprehensible than other logically equivalent forms.

The following sections define the meaning and formal properties of censored production rules and show how such rules can be used.

Censored Rules are If-Then-Unless Rules

Each rule in a production system represents a packet of knowledge that is easy to interpret, to explain, and to modify. In this paper, we write production rules in the form:

$$\begin{array}{ll} \textit{If} & \langle \textit{premise} \rangle \\ \textit{then} & \langle \textit{action} \rangle \end{array} \quad (1)$$

The $\langle \textit{premise} \rangle$ is a conjunction of predicates representing some elementary conditions, and the $\langle \textit{action} \rangle$ is what is to be done when $\langle \textit{premise} \rangle$ is satisfied.

If the $\langle \textit{premise} \rangle$ part of the rule is not satisfied, no $\langle \textit{action} \rangle$ is performed. If the $\langle \textit{action} \rangle$ part is replaced by a conjunction of predicates, then the rule becomes an implicative assertion:

$$\langle \textit{premise} \rangle \Rightarrow \langle \textit{decision} \rangle \quad (2)$$

Winston [1983] introduced the concept of a production rule augmented with an *unless* condition:

$$\begin{array}{ll} \text{If} & \langle \text{premise} \rangle \\ \text{then} & \langle \text{conclusion} \rangle \\ \text{unless} & \langle \text{censor} \rangle \end{array} \quad (3)$$

The $\langle \text{censor} \rangle$ is a logical condition (typically, a predicate or a disjunction of predicates) that, when satisfied, blocks the rule. Thus, a censor can be viewed as a statement of exceptions to the rule.

In the original formulation, given in Winston [1983], the censor is logically interpreted according to

$$\langle \text{premise} \rangle \& \neg \langle \text{censor} \rangle \Rightarrow \langle \text{decision} \rangle \quad (4)$$

which is logically equivalent to

$$\langle \text{premise} \rangle \Rightarrow \langle \text{decision} \rangle \vee \langle \text{censor} \rangle \quad (4a)$$

In this formulation, the role of the *unless* condition is similar to the definition of an *exception* described by Etherington and Reiter [1983]. The difference is that in Winston's formulation there is an additional stipulation that an unlimited effort is put into showing that $\langle \text{premise} \rangle$ is true, but only one-step effort is put into showing that $\langle \text{censor} \rangle$ is true, and when one-step effort fails, the $\langle \text{censor} \rangle$ condition is assumed to be false.

In this paper we present another interpretation of the *unless* condition, discuss its validity, and argue for the utility of rules with the new type of unless conditions.

Note that our primary intention is to develop a useful computational mechanism, rather than a cognitive model of human reasoning. Consequently the use of unless conditions in censored production rules is not intended to model of the human use of the word *unless* precisely. Nevertheless, our intuition is that censored production rules capture certain aspects of common-sense knowledge that are absent from ordinary production rules, thereby facilitating human rule creation and comprehension.

We now treat the logical aspects of unless conditions. Once that is done, we treat the expositive and control aspects.

The Unless Operator Is Logically Equivalent to Exclusive-or

Let us consider a simple statement with an *unless* condition: "If it is Sunday, John work in his yard, unless the weather is bad." Writing this statement as an if-then rule, we have the following:

<i>If</i>	it is Sunday	
<i>then</i>	John works in the yard	
<i>unless</i>	the weather is bad	(5)

If we substitute the propositional symbol S for "it is Sunday," Y for "John work in his yard," B for "the weather is bad," and denote *unless* by the symbol \backslash , then we can say S implies Y unless B , which we can write as follows:

$$S \Rightarrow Y \backslash B \quad (6)$$

Suppose we interpret (6) according to (4). Then we write this:

$$S \& \neg B \Rightarrow Y \quad (7)$$

According to (4a), we can also write this:

$$S \Rightarrow Y \vee B \quad (7a)$$

If the weather is not bad, then $\neg B$ is true. And if $\neg B$ is true and it is Sunday, then we can infer that John is in the yard. If the weather is bad, then B is true, the if-part of rule (7) is not satisfied, and nothing can be inferred about whether John is or is not in the yard.

The common-sense meaning of the expression (5), however, supports the inference that if the weather on Sunday is bad, then John does not work in the yard. Such an interpretation of expression (5) requires the following pair of assertions:

$$S \& \neg B \Rightarrow Y \quad (8)$$

$$S \& B \Rightarrow \neg Y \quad (9)$$

These assertions can be written equivalently as:

$$S \Rightarrow ((\neg B \Rightarrow Y) \& (B \Rightarrow \neg Y)) \quad (10)$$

By manipulating the *then* part, we obtain

$$S \Rightarrow ((Y \& \neg B) \vee (\neg Y \& B)) \quad (11)$$

and finally,

$$S \Rightarrow (Y \oplus B) \quad (12)$$

where \oplus denotes the exclusive-or operator.

Thus, the logical interpretation of the unless operator, \lfloor , requires it to act like the exclusive-or operator, \oplus , that connects the *then* part and the *unless* part of a censored rule.

Our unless operator takes precedence over the implication operator, \Rightarrow . This new interpretation of the *unless* operator is identical to that of the *except for* operator in variable-valued logic defined by Michalski [1980].

Comparing the *passive* definition of a censored rule given by equation (7a) and the *active* one given by equation (12), it is clear that one uses the ordinary *or* operator in the right-hand side, whereas the other uses the exclusive-or operator. These two definitions are illustrated graphically in figure 1.

The Unless Operator Makes Expectations Explicit

Censored rules are important because their unless conditions have extra-logical aspects:

- The unless operator has an *expositive* aspect because it allows us to express certain expectations.
- The unless operator has a *control* aspect because it allows us to deploy a variety of problem-solving schemes.
- The expositive and control aspects of the unless operator constitute its *pragmatics*.

In this section, we look at the unless operator from the expositive point of view. In the next, we look from the control point of view.

According to the logical interpretation, the rule

$$S \Rightarrow Y \lfloor B \quad (13)$$

is logically equivalent to

$$S \Rightarrow B \lfloor Y \quad (14)$$

and also to these negation-containing expressions:

$$S \Rightarrow \neg Y \lfloor \neg B \quad (15)$$

$$S \Rightarrow \neg B \lfloor \neg Y \quad (16)$$




Figure 1. The passive (a) and active (b) definitions of a censored rule.

Thus, the rule:

<i>If</i>	it is Sunday	
<i>then</i>	John works in the yard	
<i>unless</i>	weather is bad	(17)

should be logically equivalent to the following alternatives:

<i>If</i>	it is Sunday	
<i>then</i>	the weather is bad	
<i>unless</i>	John works in the yard	(18)

<i>If</i>	it is Sunday	
<i>then</i>	John works in the yard	
<i>unless</i>	he reads a book	(23)

Transforming according to (14), we have:

<i>If</i>	it is Sunday	
<i>then</i>	John reads a book	
<i>unless</i>	he works in the yard	(24)

Both rules seem to be reasonable and to represent logically equivalent information. But now let us represent rule (23) as a pair of rules corresponding to expressions (8) and (9):

<i>If</i>	it is Sunday and John does not read a book	
<i>then</i>	he works in the yard	(25)

<i>If</i>	it is Sunday and John reads a book	
<i>then</i>	he does not work in the yard	(26)

Rule (25) clearly implies that on Sunday John either works in the yard or reads a book. We sense that he cannot do anything else on Sunday. Also, rule (25) gives us specific information about what John does, while rule (26) seems obvious, because we know that a person cannot read a book and work in the yard simultaneously. But when a person does not read a book, he could be doing many things other than working in the yard.

Now let us represent expression (24) as a pair of rules corresponding to expressions (8) and (9):

<i>If</i>	it is Sunday and John does not work in the yard	
<i>then</i>	he reads a book	(27)

<i>If</i>	it is Sunday and John works in the yard	
<i>then</i>	he does not read a book	(28)

Here again, only the first rule supplies useful information. The second rule is redundant because we know, through common sense, that a person cannot read a book and work in the yard at the same time.

This leads us to the following observation: if the *action* and the *cursor* in a censored rule are mutually exclusive assertions, then the *active* and the *passive* interpretation of a censored rule yield the same inferences. That is, the *passive* interpretation, together with a common-sense mutual-exclusion constraint, is equivalent to the *active* interpretation.

Let us now consider another example, taken from Winston [1983].

If a man is an unmarried adult
then he is bachelor
unless he cannot be married (29)

Let us transform this rule to an equivalent one according to the same transformation that leads to (16):

If a man is an unmarried adult
then he is not a bachelor
unless he can be married (30)

Clearly, both rules, (29) and (30), are perfectly legitimate statements. Also, in both rules, the censor acts as a switch, validating or invalidating the *then* part of the rule.

There is, however, an important difference between the two rules: in one the censor condition is ordinarily false, whereas in the other, the censor condition is ordinarily true. Consequently, the following implication ordinarily holds:

If a man is an unmarried adult
then he is a bachelor (31)

Whereas the following implication ordinarily does not hold:

If a man is an unmarried adult
then he is not a bachelor (32)

Ignoring the censor in rule (29) does not invalidate the rule completely; such a truncated rule still leads to a valid conclusion most of the time. However, ignoring the censor in rule (30) does invalidate the rule almost completely. As we explain later, this leads to important expositive considerations in connection with the use of censored production rules. We will argue that the preferred form of a censored rule is the one in which the *if* and *then* parts express an important, often holding relationship, whereas the censor represents a rarely occurring exception from the rule. This helps justify Winston's assumption, mentioned earlier, that less effort should be devoted to dealing with the *censor* than to establishing the *premise*.

Summarizing, there are two important characteristics to be remembered for censored rules of the form:

$$\langle \text{premise} \rangle \Rightarrow \langle \text{decision} \rangle [\langle \text{censor} \rangle] \quad (33)$$

- First, from a logical viewpoint, the *unless* operator $[$ acts like the exclusive-or operator.
- Second, from an expositive viewpoint, the $\langle \text{censor} \rangle$ condition acts as a switch that validates or invalidates the $\langle \text{decision} \rangle$. Its desired use

is when the assertion $\langle \text{premise} \rangle \& \langle \text{decision} \rangle$ holds frequently and the assertion $\langle \text{premise} \rangle \& \langle \text{censor} \rangle$ holds rarely.

The Unless Operator Supports Various Control Schemes

So far, we have discussed the unless operator from two points of view: from the logical point of view, we showed that it is equivalent to the exclusive-or operator; and from the expositive point of view, we showed that it makes certain expectations explicit.

Because expectations are made explicit, there are a number of obvious alternate control schemes for using censored rules. Here are the two extreme possibilities:

- The Missouri method: Treat the unless operators as if they were exclusive-or operators, thereby ignoring expectation information. Good in situations for which expectations are unreliable and for which nothing should be assumed.
- The ask-questions-later method: Ignore all censors. Good in situations for which rapid response is critical.

And of course, there are intermediate possibilities, of which the following are representative:

- The decision-maker method: Assume all censor conditions are false unless already known to be true. Put no effort into showing that a censor condition is true.
- The trusting-skeptic method: Once a premise is established, try to show that censor conditions are true, but fix the depth of rule chaining to some prescribed number of levels or to some prescribed consumption of resources. We call this the trusting-skeptic method because it reflects an assumption that the indicated expectations are solid and that there is little point in putting more than a little effort into overturning those expectations. This is the original method proposed by Winston.
- The stubborn-donkey method: Do not allow situations in which censors are triggered by rules that themselves have censors that are triggered by other rules ad nauseum. Fix the depth of censor chaining to some prescribed number of levels or to some prescribed consumption of resources. We call this the stubborn-donkey method because the number of times a decision can be reversed is limited.
- The tapered-search method: Allow any number of levels of censor chaining, but reduce the resources allocated to showing that censors are true in proportion to the depth of chaining.

Importantly, for all these and all similar schemes, tentative answers can be reported as soon as the premise part of the rule allows. Then the censor parts of the rules can be pursued according to whatever resource allocation scheme seems best. At any moment, the tentative answer is the best answer possible relative to the chosen resource allocation scheme and the expended resources. In principle, tentative answers may change many times, but those tentative answers probably will not change much in practice.

From this point of view, our control proposals are reminiscent of the progressive-deepening notion developed for chess-playing programs [Frey, 1983] [Winston, 1984]. Progressive deepening enables game-playing programs to produce reasonable moves quickly, with better moves forthcoming if the clock allows.

From another point of view, our proposals are reminiscent of the work of Carl E. Hewitt, whose thesis argued persuasively that logical rules should be augmented with knowledge about appropriate uses [Hewitt, 1972]. In the thesis, the primary distinction was between antecedent and consequent rules.

While on the subject of control, two other questions spring to mind:

- Why not devise a parallel-processing scheme?
- Why not devise a more refined, quantitative representation of expectation that would support more sophisticated control schemes?

These ideas seem natural to use, but we have not yet thought much about parallel-processing, nor have we devised control schemes exploring quantitative representation. We have developed a quantitative representation, however, which we describe in the next section.

The Augmented Unless Operator Makes Expectations Quantitative

Let us now give a more quantitative definition of a censored rule. Consider the following rule:

$$P \Rightarrow D[C] \tag{34}$$

where P is a premise, D is a decision, and C is a censor. Although the unless operator, $[$, is logically equivalent to the commutative exclusive-or operator, the unless operator has an expositive aspect which is not commutative. In order to capture the asymmetry precisely, let us associate two parameters, γ_1 and γ_2 , with rule (34).

$$P \Rightarrow D[C : \gamma_1, \gamma_2] \tag{31}$$

Both γ_1 and γ_2 are point probabilities, one indicating the strength of the relationship between P and D , and the other, between P and C .² Now consider the following sets:

- Ω is a finite sample of events.
- Ω_P is the set of events for which P holds.
- Ω_{PD} is the subset of events for which both P and D hold.
- Ω_{PC} is a subset of events for which both P and C hold.

Given these sets, the parameters γ_1 and γ_2 are defined as follows:

$$\gamma_1 = \frac{\Pr[P, D]}{\Pr[P]} = \Pr[D|P] \simeq \frac{|\Omega_{PD}|}{|\Omega_P|} \quad (32)$$

$$\gamma_2 = \frac{\Pr[P, C]}{\Pr[P]} = \Pr[C|P] \simeq \frac{|\Omega_{PC}|}{|\Omega_P|} \quad (33)$$

where $|\Omega_i|$ denotes the cardinality of Ω_i . Also we assume that $\Omega_D \cap \Omega_C = \emptyset$ and $\Omega_D \cup \Omega_C = \Omega_P$; thus $\Pr[D|P] + \Pr[C|P] = 1$.

Relating these definitions to our example about John working in the yard unless weather is bad, we can say that Ω is a set of days over a sufficiently large period of time; Ω_P is the set of Sundays during this period of time; Ω_{PD} is the set of Sundays when John works in the yard; and Ω_{PC} is the set of Sundays with bad weather.

Assuming that there are significantly more Sundays when John works in the yard than there are Sundays when the weather is bad, then Ω_{PD} is considerably larger than the set Ω_{PC} :

$$|\Omega_{PD}| \gg |\Omega_{PC}| \quad (34)$$

Thus, taking into consideration (32) and (33), we have

$$\gamma_1 \gg \gamma_2 \quad (35)$$

In our example, γ_1 stands for the ratio of Sundays when John worked in the yard to all Sundays, and γ_2 stands for the ratio of Sundays with bad weather to all Sundays.

And from the logical point of view, according to our interpretation of the unless operator, the sets Ω_{PD} and Ω_{PC} must be disjoint. Consequently, the sum $\gamma_1 + \gamma_2$, must always equal 1. Thus, knowing γ_1 , it is easy to compute γ_2 ; therefore rule (31) can be simplified:

² It may be better to introduce parameters indicating belief and disbelief for both D and C , as in MYCIN, or to introduce parameters indicating lower bounds on the probabilities of the truth and the falsehood of both D and C [Shafer, 1976], as developed in Quinlan's INFERNO system [1983]. For simplicity, we consider here only subjective point probabilities.

$$P \Rightarrow D|C : \gamma \quad (36)$$

where γ stands for γ_1 and $\gamma \geq 0.5$. Because of (35), $\gamma = \gamma_1$ should be significantly greater than 0.5, but we will assume only $\gamma \geq 0.5$. If $\gamma = 1$, then $P \& C$ never holds, and rule (36) becomes

$$P \Rightarrow D$$

Note that γ does not say anything about how often relations $\neg P \& D$ and $\neg P \& \neg D$ hold.

Now suppose that γ is 0.9 in the following rule:

$$P \Rightarrow D|C : \gamma \quad (37)$$

If P holds, rule (37) allows us to make inference that D and $\neg C$ hold with certainty 0.9 and $\neg D$ and C will certainly 0.1. Thus if we know P holds and do not know if C holds, we infer that D holds with certainty 0.9. On the other hand, if P holds and C does not, then we can infer that D holds with certainty 1.

If both P and C hold, then we can infer that D does not hold. Symmetrically, if both P and D hold, we can infer that C does not hold.

Now suppose that we know whether D holds but do not know whether P or C hold. If we ignore γ , then (37) can be written as a pair of expressions:

$$P \& \neg C \Rightarrow D \quad (38)$$

$$P \& C \Rightarrow \neg D \quad (39)$$

Expressions (38) and (39) can be rewritten as follows:

$$\neg D \Rightarrow \neg P \vee C \quad (40)$$

$$D \Rightarrow \neg P \vee \neg C \quad (41)$$

If $\neg D$ holds, then either $\neg P$ or C . If we know that $\neg C$ holds, then using (40) we can infer with certainty 1.0 that $\neg P$. If C , then nothing can be said about $\neg P$ in this case.

Similarly, if D holds and we know that C holds, then we can infer with certainty 1.0 that $\neg P$. If $\neg C$, then nothing can be said about $\neg P$.

Thus rule (37) permits us to generate a number of inferences of varying certainty, depending on what is given and what is unknown. Also, there is a natural relationship between the certainty of decisions and the amount of knowledge available.

The relationship between certainty and knowledge has an important operational consequence. To illustrate, consider two cases:

- P is known to hold, and there are insufficient time or space to determine C . A system can infer the decision D , with certainty γ .
- P is known to hold, and there are sufficient resources to determine C . A system can determine C and subsequently conclude D or $\neg D$, depending on C , with certainty 1.

Rules May Have Many Censors

A censored production rule may have more than one exception-denoting censor. Consider, for example, the assertion that birds fly:

$$\forall x \text{ is-bird}(x) \Rightarrow \text{flies}(x) \quad (42)$$

This general assertion enables us to expect that any newly observed bird flies. But not all birds fly. For example, penguins, ostriches, emus, kiwis, and domestic turkeys do not fly. To include this information, we write:

$$\begin{aligned} \forall x \text{ is-bird}(x) \Rightarrow \text{flies}(x) \mid & (\text{is-penguin}(x) \\ & \vee \text{is-ostrich}(x) \\ & \vee \text{is-emu}(x) \\ & \vee \text{is-kiwi}(x) \\ & \vee \text{is-domestic-turkey}(x)) \end{aligned} \quad (43)$$

Thus the exceptions are disjunctively linked together as one censor condition. Suppose we generalize these exceptions into one statement for special birds. Then we can write:

$$\forall x \text{ is-bird}(x) \Rightarrow \text{flies}(x) \mid \text{is-special-bird}(x) \quad (44)$$

But then, the rule is still not entirely correct. Even a flying bird cannot fly when it is dead or sick or has broken wings. Let us characterize all these situations as bird being in an unusual condition. Then we can write:

$$\begin{aligned} \forall x \text{ is-bird}(x) \Rightarrow \text{flies}(x) \mid & (\text{is-special-bird}(x) \\ & \vee \text{is-in-unusual-condition}(x)) \end{aligned} \quad (45)$$

where

$$\begin{aligned} \forall x \text{ is-special-bird}(x) \Leftarrow & \text{is-penguin}(x) & (46) \\ & \vee \text{is-ostrich}(x) \\ & \vee \text{is-emu}(x) \\ & \vee \text{is-kiwi}(x) \\ & \vee \text{is-domestic-turkey}(x) \end{aligned}$$

and

$$\begin{aligned} \forall x \text{ is-in-unusual-condition}(x) \Leftarrow & \text{is-dead}(x) & (47) \\ & \vee \text{is-sick}(x) \\ & \vee \text{has-broken-wings}(x) \end{aligned}$$

Now a bird also cannot fly when its legs are stuck in concrete. This case may also be classified as bird in an unusual condition. Thus, to update our knowledge, we need not change our basic rule; we need only extend our definition of unusual condition:

$$\begin{aligned} \text{is-in-unusual-condition}(x) \Leftarrow & \text{is-dead}(x) & (48) \\ & \vee \text{is-sick}(x) \\ & \vee \text{has-broken-wings}(x) \\ & \vee \text{has-legs-stuck-in-concrete}(x) \end{aligned}$$

Thus censored production rules facilitate small rule repairs.

Rules May Have Incomplete Censors

Censors are generalized whenever a disjunctive condition is added, as in going from (44) to (45) or from (47) to (48). The generalized censor, (48), fires in additional situations, preventing the rule from asserting the decision. Thus, generalizing a censor specializes a censored rule. Conversely, specializing the censor (up to its complete removal) generalizes a censored rule.

Let us now go back to rule (42) and augment it by adding the parameter γ :

$$\forall x \text{ is-bird}(x) \Rightarrow \text{flies}(x) \lfloor \text{is-special-bird}(x) : \gamma \quad (49)$$

where γ estimates the probability that any given bird flies. In the case that a given bird does not fly, then according to rule (49) the censor

is-special-bird(x) must be true. This is at odds with (45), however, because the bird may be not be a special bird, but rather, in an unusual condition.

Instead of generalizing the censor to correct rule (49), let us introduce a parameter δ that adds an additional uncertainty to compensate for the incompleteness of the censor.

$$\forall x \text{ is-bird}(x) \Rightarrow \text{flies}(x) \mid \text{is-special-bird}(x) : \gamma, \delta \quad (50)$$

Here is the intended meaning:

- γ is the certainty that a bird flies when we do not know whether it is a special bird or not.
- δ is the certainty that a bird flies when we know that it is not a special bird.

Thus δ accounts for the fact that the censor is incomplete.

Now let us discuss the meaning of parameters γ and δ generally. Suppose we insert into a censored rule a symbol UNK standing for *UNKNOWN*:

$$P \Rightarrow D \mid (C_1 \vee \text{UNK}) : \gamma, \delta \quad (51)$$

The symbol UNK represents a disjunction of unknown conditions that could block the inference of D from P .

If we know that the condition C_1 does not hold ($\neg C_1$), then the strength of the implication $P \Rightarrow D$ depends on the condition UNK. If none of the UNK conditions hold, then D holds. If any of the UNK conditions hold, then D does not.

Parameter δ is defined to be the certainty that $P \Rightarrow D$ when $\neg C_1$ is true, that is $\delta = \text{Pr}[D \mid P, \neg C_1]$. This is equivalent to the certainty that there is no implicit part of the censor that can hold when $\neg C_1$ is true.

Parameter γ is defined as the certainty that $P \Rightarrow D$ when it is not known whether $(C_1 \vee \text{UNK})$ holds, that is $\gamma = \text{Pr}[D \mid P]$. The implication $P \Rightarrow D$ has certainty 1 when $C_1 \vee \text{UNK}$ is known to be false. Thus, the parameter γ is equivalent to the *a priori* certainty that none of the censors hold.

Obviously the *a priori* certainty of $\neg(C_1 \vee \text{UNK})$ must be equal to or smaller than the *a priori* certainty that $\neg \text{UNK}$. Therefore, $\gamma \leq \delta$. Note that $\delta = 1$ if it is certain that there are no conditions in the censor other than C_1 .

Let us rewrite rule (51) as two rules:

$$P \Rightarrow D \mid C^* : \gamma, \delta \quad (52)$$

$$C^* \Leftrightarrow C_1 \vee \text{UNK} \quad (53)$$

Expression (53) can be rewritten as

$$\neg C^*[(C_1 \vee \text{UNK})] \quad (54)$$

In this form it clearly states that the censor will not fire unless C_1 is true or UNK is true. Suppose now that P holds and we ignore censor C^* , that is, we ignore C_1 and UNK, then we can conclude D with the certainty γ . Parameter γ is therefore called the 0-level strength of implication $P \Rightarrow D$ (because no information about the censors is taken into account). If we evaluate C_1 , and C_1 is false, then from P we can conclude D with the certainty δ . Parameter δ is thus called the 1-level strength of implication. Recapitulating our discussion, we define a censored rule as follows:

$$P \Rightarrow D[C : \gamma, \delta] \quad (55)$$

where P is the premise, D is the decision, C is the censor,

γ is the *a priori* certainty that $P \Rightarrow D$ when C is unknown (the 0-level strength of implication), δ is the *a priori* certainty that $P \Rightarrow D$ when C is known to be false (the 1-level strength of implication).

The Provided Operator Complements the Unless Operator

Consider this statement:

<i>If</i>	It is Saturday	
<i>then</i>	I will go to a concert	
<i>unless</i>	I cannot get a babysitter	(56)

We can express this in our logic as follows:

$$S \Rightarrow C[B] \quad (57)$$

where S is for *on Saturday*; C is for *I will go to a concert*; and B is for *I cannot get a babysitter*.

An alternate way to say the same thing is as follows:

<i>If</i>	It is Saturday	
<i>then</i>	I will go to a concert	
<i>provided</i>	I can get a babysitter	(58)

Similarly, statements such as the following can be expressed using the *provided* form.

$$\text{You will enjoy the hike unless the weather is bad} \quad (59)$$

$$\text{A happy future is guaranteed unless we do not have peace} \quad (60)$$

Changing the polarity of the *cursor* condition replaces *unless* by *provided*; a positive cursor used with an *unless* condition becomes a negative cursor with the *provided* condition and vice versa.

Of course, in normal human use, the word *provided* introduces a pre-condition rather than an exception.

Let us now introduce the *provided* operator, denoted \lceil , to complement the *unless* operator. Thus, the rule

$$P \Rightarrow D \lfloor C \quad (61)$$

is assumed to be equivalent to

$$P \Rightarrow D \lceil (-C) \quad (62)$$

and the rule

$$P \Rightarrow D \lfloor (-C) \quad (63)$$

is assumed to be equivalent to

$$P \Rightarrow D \lceil C \quad (64)$$

From the logical standpoint, the exclusive-or operator in $D \oplus \neg C$ is equivalent to $D \equiv C$. Thus, from the logical viewpoint, the *provided* operator acts as the *equivalence* operator.

Because *unless* and *provided* operators are complementary, any censored rule can be expressed using only a positive cursor condition. Also, replacing one operator by another does not effect the strength of implication parameters γ and δ :

$$P \Rightarrow D \lfloor C : \gamma, \delta \quad (65)$$

guarantees

$$P \Rightarrow D \lceil C' : \gamma, \delta \quad (66)$$

where $C' = \neg C$.

There is a small difference in the rule interpretation: in rule (66), δ is interpreted as the strength of implication $P \Rightarrow D$, when C' holds. The parameter γ remains to denote the strength of implication when we do not know if C' holds or not.

Representing the In-which-case Condition

Suppose we want to represent the statement, "On Sunday I will fly a kite, *unless* there is no wind, *in which case* I will write poetry."

The *unless* operator does not enable us to express this statement directly. We could introduce an *in-which-case* operator, but we prefer to use two statements instead: "On Sunday I will fly a kite *provided* there is wind," and "On Sunday I will write poetry *provided* there is no wind." These statements are directly expressible as censored production rules:

$$S \Rightarrow K|W \quad (67)$$

$$S \Rightarrow P|-W \quad (68)$$

Note that (68) can be expressed as a rule with an *unless* condition:

$$S \Rightarrow P|W \quad (69)$$

The English form sounds strange, however, in light of (67): "On Sunday I will write poetry *unless* there is wind". Rule (68), with a negative censor, seems more appealing than rule (69) with a positive censor. There seems to be a general regularity: from the expositive viewpoint, it is better to use complementary censor conditions and the same censor operator, rather than to use complementary operators and the same censor condition.

Why use censored rules rather than ordinary production rules:

$$S \& W \Rightarrow K \quad (70)$$

$$S \& \neg W \Rightarrow P \quad (71)$$

The answer depends on what we want to express. Rule pairs (67, 68) and (70, 71) are not exactly equivalent. The first pair, (67, 68), can be re-expressed as one expression:

$$S \Rightarrow K|W \& P|-W \quad (72)$$

This in turn can be rewritten as

$$S \Rightarrow W \& K \vee \neg W \& P \vee \neg(K \vee P) \quad (73)$$

The second rule pair, (70, 71), can be re-expressed as

$$S \Rightarrow (W \Rightarrow K) \& (\neg W \Rightarrow P) \quad (74)$$

and then as

$$S \Rightarrow W \& K \vee \neg W \& P \vee K \& P \quad (75)$$

Thus, the difference between the two pairs of rules, (67, 69) and (70, 71), is in the third component of expressions (73) and (75). The first pair implies that a person may neither fly a kite nor write poetry on Sunday, as well as that he will not simultaneously fly a kite and write poetry. The second pair implies that he must do one thing or the other and accepts the possibility that both may be done at once. (This is not to say that flying a kite and writing poetry are mutually exclusive, but only that there is a difference between these two pairs of rules.)

Distinguishing Rules from Definitions

Consider the following expression, which represents the sentence, "On Sunday John will fly a kite *provided* there is wind:"

$$S \Rightarrow K | W \quad (76)$$

Suppose we also want to express some knowledge about whether there will be wind on Sunday. Assume that this condition is: "There will be wind on Sunday (*W*) if there is drop in the temperature on Friday (*DF*) *unless* Saturday is sunny (*SS*). We can write this condition as a production rule:

$$DF \Rightarrow W | SS \quad (77)$$

Rule (76) is typically used in the forward direction for answering the question "What will John do on Sunday?" On the other hand, rule (77) is typically used in the backward direction for answering the question "Will there be wind on Sunday?" Rules evoked in the backward direction are definitions. To reflect the differences in control of rule execution, as well to facilitate human readability, it is desirable to make a distinction in the form of a forward-executed rule and a backward-executed definition (though logically they are equivalent). Such a distinction can be made very simply by writing the rule (77) in the following form:

$$W \Leftarrow DF | SS \quad (78)$$

This is to be read as *W* if *DF* unless *SS*. The \Leftarrow is interpreted as a logical implication, but directed in the opposite direction than normal. We used this form already in expressions (46) and (47), but without the *unless* condition. Note that in this case, in order to interpret the rule correctly, the censor *SS* is assumed to be linked by the exclusive-or operator with the *left side* of \Leftarrow rather than with the *right side*.

Also note that the operator linking a term being defined with the body of the definition is typically equivalence. Thus, to have a complete logical representation of a definition, one should use equivalence rather than implication.

Inference Rules for Transforming Censored Production Rules

This section gives a sample of inference rules applicable to censored production rules. In these inference rules, the premise and the decision involve rules. A decision rule is a logical consequence of premise rules. To state that an assertion A has truth value α we write

$$A : \alpha \quad (79)$$

To express an inference rule stating that A is a logical consequence of A_1, A_2, \dots , we write

$$\begin{array}{l} A_1 : \alpha_1 \\ A_2 : \alpha_2 \\ \vdots \end{array} \left| \right. > A : \alpha \quad (80)$$

We will drop the certainty parameters δ and γ in censored rules whenever they are irrelevant for the given inference rule.

$$\begin{array}{l} P \Rightarrow D[C : \gamma, \delta \\ P : \text{True} \\ C : \text{Unknown} \end{array} \left| \right. > D : \gamma \quad (81)$$

$$\begin{array}{l} P \Rightarrow D[C : \gamma, \delta \\ C : \text{False} \end{array} \left| \right. > P \Rightarrow D : \gamma \quad (82)$$

$$\begin{array}{l} P \Rightarrow D[C : \gamma, \delta \\ D : \text{True} \end{array} \left| \right. > P \Rightarrow \neg C : \text{True} \quad (83)$$

$$\begin{array}{l} P_1 \Rightarrow D[C \\ P_2 \Rightarrow D[\end{array} \left| \right. > P_1 \vee P_2 \Rightarrow D[C \quad (84)$$

$$\begin{array}{l} P \Rightarrow D_1[C \\ P \Rightarrow D_2[C \end{array} \left| \right. > P \Rightarrow D_1 \& D_2[C \quad (85)$$

$$\begin{array}{l} P \Rightarrow D[C_1 \\ P \Rightarrow D[C_2 \end{array} \left| \right. > P \Rightarrow D[(C_1 \vee C_2) \quad (86)$$

$$\begin{array}{l} P \Rightarrow D | C_1 \\ C_1 | C_2 \end{array} \Bigg| > P \Rightarrow D | C_2 \quad (87)$$

$$\begin{array}{l} P \Rightarrow D | C_1 \\ D \Rightarrow D_1 | C_2 \\ D_1 \Rightarrow D \end{array} \Bigg| > P \Rightarrow D_1 | (C_1 \vee C_2) \quad (88)$$

In rule (88), the implication $D_1 \Rightarrow D$ is needed to prevent the possibility of having C_1 and D_1 hold simultaneously.

For illustration of (88), let us consider an example. Suppose that the input assertions are:

On Sundays John Brite goes to a park unless he is writing poems (89)

In the park he flies a kite unless there is no wind (90)

He flies a kite only in the park (91)

The inference rule (88) allows us to make the following deduction:

On Sundays John Brite flies a kite (92)
unless he is writing poems or there is no wind

Variable Precision Logic is Related to Non-Monotonic Logic

As mentioned earlier, decisions from censored production rules depend on the truth value of the censors. The censors are assumed to be low-likelihood assertions. Therefore, if there is insufficient time or other resources to determine the the censors' truth values, the censors can be ignored, and the decision is still true with high likelihood. These decisions may have to be revised, if the censors are found later to hold. Thus, the mechanism of censored production rules enables one to make revisions in once-accepted decisions.

This is reminiscent of non-monotonic logic, which specifically investigates problems of revising beliefs and modifying tentative knowledge. In non-monotonic logic, the basic rule of inference is "If a negation of a formula is not derivable from axioms by inference rules of the first order theory, then accept the formula as true." To formalize such a rule, this logic extends the classical logic by introducing a proposition-forming modality M . A proposition Mp states that p is consistent with everything believed. A

comprehensive treatment of various theoretical aspects of non-monotonic logic, such as consistency and provability is given by McDermott and Doyle [1980].

A related concept, circumscription, was introduced by McCarthy [1980]. In reasoning by circumscription, one assumes that the only objects that have some property are those that can be shown to have it.

The approach taken in non-monotonic reasoning is quite different from the one taken in variable precision logic. We do not introduce any new modalities, but rather a new operator (the unless operator). Also, unlike statements in non-monotonic logic, censored production rules can be associated with certainty parameters and estimates of the cost of pursuing individual conditions. These certainty parameters and costs can be used for controlling the use of the rules and estimating the certainty of decisions in any act of inference. Thus, the variable precision logic adopts some aspects of both non-monotonic logic and many-valued logic.

Summarizing, the goals and methods of variable precision logic are different from those of non-monotonic logic. While non-monotonic logic investigates the formal implications from reasoning involving uncertain assumptions, variable precision logic attempts to develop mechanisms for representing and conducting reasoning that reflects different trade-offs between the certainty (and/or specificity) of decisions and the computational resources needed to derive them.

Conclusion

Classical logic was conceived originally as a prescriptive theory of how an ideal mind might reason. For any inference, it assumes that all needed premises are known in advance and that the truth values of those premises do not change. It ignores the time and memory resources needed for reasoning. In the real world, however, both humans and computers often must reason using insufficient, incomplete, or tentative premises. Moreover, both are subject to constraining time and memory limitations.

Nevertheless, both humans and computers must be able to react promptly to new information, and they must be able to change or repair their knowledge when new information produces contradictions or when initial assumptions are withdrawn.

Here we have focused on one aspect of the problem by describing a simple knowledge representation and a reasoning system that enables trade-offs between the certainties of various decisions and the effort needed to derive those decisions.

We showed that by factoring out conditions that have low likelihood and by treating them outside of the main line of reasoning, a system can easily exercise different control schemes for rule execution.

Common-sense reasoning seems to follow the most important and likely lines of argument, ignoring myriad low-likelihood exceptions. Variable-precision logic, through the mechanism of censored production rules, provides a simple computational mechanism for capturing some of the properties of such reasoning. The same censored-production-rule mechanism also facilitates minor repairs to the rules.

We also suggest that a computationally-limited reasoning system should associate rule premises and censors with numerical estimates of their likelihood and testing cost. These estimates enable an inference system to decide which premises and censors to evaluate under given time and cost constraints.

The control planning problem is one of several important topics that were not discussed here, but which clearly invite further research. For example, we need to understand how a reasoning mechanism can make controllable trade-offs between certainty and specificity.

We also need to devise efficient algorithms for learning censored production rules from specific cases and precedents and to modify those rules incrementally to account for new facts. Some initial work in this direction was recently done by Becker [1985].

Acknowledgments

The authors thank Jeff Becker, Boris Katz, Karen Prendergast, and the anonymous reviewers of the *Artificial Intelligence Journal* for thoughtful and detailed comments, which helped us to improve this paper.

References

- Becker, J. M., "Inductive Learning of Decision Rules with Exceptions: Methodology and Experimentation," MS Thesis, Department of Computer Science, University of Illinois, Urbana, IL, July, 1985.
- Etherington, David W. and Raymond Reiter, "On Inheritance Hierarchies with Exceptions," *Proceedings of the National Conference on Artificial Intelligence*, August 22-26, 1983.
- Hewitt, Carl E., "Description and Theoretical Analysis (using schemata) of PLANNER: a Language for Proving Theorems and Manipulating Models in a Robot," PhD Thesis, MIT AI Laboratory, TR-258, 1972.
- McDermott, Drew and Johnathan Doyle, "Non-Monotonic Logic I," *Artificial Intelligence*, volume 13, 1980.
- Frey, Peter W., *Chess Skill in Man and Machine*, Springer-Verlag, 1983.
- McCarthy, John, "Circumscription—A Form of Non-Monotonic Reasoning," *Artificial Intelligence*, volume 13, 1980.
- Michalski, R. S., "Pattern Recognition as Rule-Guided Inductive Inference," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-2, No. 4, July 1980, pp. 349-361.

- Michalski, Ryszard S., "A Theory and Methodology of Inductive Learning," in *Machine Learning—An Artificial Intelligence Approach*, edited by Ryszard S. Michalski, Jaime G. Carbonell, and Tom M. Mitchell, Tioga Publishing Co., Palo Alto, CA, 1983.
- Minsky, Marvin, "The Society of Mind", forthcoming.
- Reinke, R. and Ryszard S. Michalski, "Incremental Learning of Concept Descriptions, *Machine Intelligence 11*, J. E. Hayes, D. Michie, and J. Richards, eds., Oxford University Press, 1985.
- Shafer, G., *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, 1976.
- Quinlan, J. Ross, "Inferno: A Cautious Approach to Uncertain Inference," *The Computer Journal*, volume 26, number 3, 1983.
- Winston, Patrick Henry, *Artificial Intelligence, Second Edition*, Addison-Wesley, 1984.
- Winston, Patrick Henry, "Learning New Principles from Precedents and Exercises," *Artificial Intelligence*, volume 19, number 3, 1981.
- Winston, Patrick Henry, "Learning by Augmenting Rules and Accumulating Censors," MIT AI Laboratory, AIM-678, 1982. A version appears in *Machine Learning: An Artificial Intelligence Approach Volume II*, edited by Ryszard S. Michalski, Jaime G. Carbonell, and Tom M. Mitchell, 1986.

