An Endorsement-Based Plan Recognition Program

Michael Sullivan and Paul R. Cohen
Department of Computer and Information Science
University of Massachusetts, Amherst
COINS Technical Report 84-33

EKSL Report 2

This research is funded by a grant from the National Science Foundation IST 8409623.
Published in *Proceedings of IJCAI-85*.

An Endorsement-Based Plan Recognition Program

Michael Sullivan and Paul R. Cohen

COINS Technical Report 84-33

December, 1984

Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01003

This research is funded by a grant from the National Science Foundation, number IST-8409623.

Abstract

This paper describes a plan-recognition program built to explore the theory of endorsements (Cohen, 1983). The program evaluates alternative interpretations of user actions and reasons about which are the most likely explanation of the user's intentions. Uncertainty about the various alternatives was encoded in data structures called *endorsements*. The paper describes the workings of this program and the successes and limitations of the endorsement-based approach.

1.0 Introduction

Historically, expert systems and other artificial intelligence programs have employed numerical techniques to assess subjective degrees of belief in uncertain alternatives. The advantages of this approach are apparent: alternatives are easily ranked by their degrees of belief, and the degrees of belief in combinations of evidence are easily calculated by simple arithmetic rules. The disadvantages are also apparent. Subjective degrees of belief generally do not behave as probabilities (Tversky and Kahneman, 1974), so we must wonder what interpretation to give them, and whether probabilistic combining methods are valid. Expert system builders report that experts are uncomfortable committing themselves to numbers (e.g., Gadsden, 1984); the numbers may be ambiguous – composed of salience and probability considerations (Buchanan and Shortliffe, 1984, Chapter 10); and, depending on how the numbers are used, their accuracy has little or no effect on performance (Buchanan and Shortliffe, 1984, Chapter 10). If the interpretations of subjective degrees of belief are unclear (e.g., if they are ambiguous) then combining degrees of belief only clouds their semantics further. Finally, numbers represent very little about uncertainty: they tell us *how much* to believe, not *why* to believe.

The theory of endorsements answers some of these concerns about numbers. Endorsements are records of sources of uncertainty; they provide explicit records of the introduction of uncertainty into reasoning. This paper describes a program that reasons explicitly about uncertainty in plan recognition problems. The first section introduces the theory of endorsements. The next section describes a plan recognition program that uses endorsements to reason about the uncertainty in interpreting user actions. We then discuss some of the problems we encountered in implementing this program, problems which highlight some of the key issues involved in dealing with endorsements. The final section is a discussion and summary.

2.0 Endorsements

Endorsements represent knowledge about uncertain situations (Cohen, 1983) This knowledge can include, but is not limited to, reasons to believe and disbelieve uncertain propositions. The major advantage of endorsements is that they make sources of uncertainty *explicit*, so we may reason about them directly, instead of implicitly through some sort of numerical calculus.

Although we give endorsements evocative, mnemonic names in this paper (e.g., "could be a mistake"), they are just tokens. Their meanings derive from specifying the situations in which they are applicable, how they combine, and how they are ranked. To build an endorsement-based system, then, one first identifies and names the sources of certainty and uncertainty in a domain. The names are called endorsements. Then, one specifies how these sources interact, so that more or less certain evidence can be combined. Finally, one gives rules for ranking combinations of sources of certainty and uncertainty, so that decisions can be made. Each of the three specifications will be discussed in the context of an example.

3.0 HMMM - An Endorsement-Based Plan Recognition Program

HMMM is a plan recognition program that infers which of several known plans a user intends by combining the evidence provided by successive user actions. Plan recognition is uncertain for two reasons: the user might make a mistake, in which case extrapolating from the action might suggest the wrong plan; or a user action may be ambiguous; i.e., the action might be consistent with several known plans.[1] If all the user's actions belong to only one known plan, the interpretation process is straightforward; but when an action can be interpreted as a mistake, or as belonging to more than one plan, HMMM is uncertain of the user's intentions, and so generates endorsements for the competing interpretations. HMMM is a simplified version of POISE (Carver, Lesser, and McCue, 1984), an office automation system with an intelligent user interface, which discerns a user's plan and offers assistance by automating some plan steps.

Individual plan steps are interpreted in the context of developing plans. The program uses its knowledge of the user's previous actions to restrict the interpretations of the current action. For example, assume the program knows the following plans:

---

[1] Other sources of uncertainty in plan recognition include an incomplete library of known plans and inaccuracies in the plan library. We limited our exploration to unintended and ambiguous actions.

Plan	Steps
plan1	a b d
plan2	b d e
plan3	a c d

Given that the user takes the actions **a** followed by **b**, we can construct three interpretations for each action:

(start plan1 **a**)	(continue plan1 **b**)
(start plan3 **a**)	(start plan2 **b**)
(mistake **a**)	(mistake **b**)

However, the interpretation of **b** as continuing plan1 would not be valid unless the first step of plan1, **a**, had already been taken. We account for these syntactic restrictions with data structures called *step linkages* . Each step linkage represents an interpretation of all the plan steps taken so far. Step linkages for the "current" step are constructed from the existing step linkages, which link all previous steps. For an interpretation that continues an already-opened plan (as **b** above continues plan1), each step linkage that mentions the preceding step is extended to include the new step. For an interpretation of a plan step as starting a new plan (as **b** above is interpreted as starting plan2), *all* step linkages are extended to include this interpretation.

Each step linkage carries endorsements. These are reasons to believe and disbelieve the interpretations of plan steps represented by the step linkages. For example, a reason to believe that **b** continues plan1, above, is that "continuity is desirable." Recall our contention that these reasons have no implicit meaning, no matter how evocative are the strings we use. The following example shows how meaning is ascribed to endorsements and how endorsements facilitate reasoning about uncertain interpretations.

4.0 An Example of Endorsement-Based Plan Recognition

Suppose we have a simple environment in which we know that the user intends exactly one of two known plans,

Plan	Steps
plan1	a b c
plan2	b d e

and the user types the input actions **a** followed by **b** followed by **d**. Briefly, we can imagine interpreting the first input as evidence for plan1, and the second as further evidence. The third input lends support for the plan2 interpretation of **b**, and casts doubt on the plan1 interpretation of a, and indirectly supports the possibility that a was a mistake. If a fourth input was c, we'd want the system to reaffirm its belief in plan1, whereas an input of e should have the opposite effect.

4.1 Applicability Conditions for Endorsements

HMMM uses endorsements to reason as just described. The actions a, b, d result in the following syntactic interpretations:

Step	Interpretation	Endorsements
1: a	(start plan1 a)	(a only grammatical possibility +) (a could be a mistake -)
2: b	(continue plan1 b)	(a b continuity is desirable +) (b other grammatical possibility -) (b could be a mistake -)
b	(start plan2 b)	(a b discontinuity is undesirable -) (b other grammatical possibility -) (b could be a mistake -)
3: d	(continue plan2 d)	(d only grammatical possibility +) (b d continuity is desirable +) (d could be a mistake -)

The endorsements are associated with the interpretations by rules specifying their applicability conditions: "other grammatical possibility" is applicable whenever a plan step figures in more than one possible plan; "could be a mistake" is always applicable; "continuity is desirable" is redundant with the interpretation of a plan step as continuing an open plan; and "discontinuity is undesirable" applies whenever a plan step is interpreted as disrupting an already open plan by starting a new one. Some endorsements are *positive*, meaning that they support the interpretation with which they are associated. Others are *negative* – reasons to disbelieve their associated interpretations.[2]

---

[2] Applicability conditions for endorsements include rules to decide whether an endorsement is positive or negative. This is easy in HMMM, but we believe it to be difficult in general to decide whether evidence speaks for or against a hypothesis.

4.2 Combining Endorsements

The endorsements associated with an interpretation are brought along when that interpretation is appended to a step linkage, and they are combined with endorsements from the previous steps in the linkage to give the endorsements of the plan up to that point. For example, the input **a** is evidence for plan1, and **b** is *further* evidence for plan1. Note that **b** is a different kind of evidence from **a**, because it is ambiguous between plan1 and plan2. Applicability conditions for endorsements give us the mechanism to distinguish between the kinds of evidence – each kind carries characteristic endorsements – but they don't specify how to *combine* the endorsements of pieces of evidence, such as **a** and **b**, when they support the same hypothesis (in this case, plan1). To this end, we have implemented *semantic combining rules*, two of which follow.

SCR1: If (plan N: step i could be a mistake -) and
(plan N: steps i j continuity is desirable +)
Then erase (plan N: step i could be a mistake -)

SCR2: If (plan M: steps i j discontinuity is undesirable -) and
(plan M: steps j k continuity is desirable +) and
(plans N,M: step j other grammatical possibility -)
Then erase (plan M: steps i j discontinuity is undesirable -)

Both rules use the occurrence of two consecutive plan steps as a basis for removing negative endorsements that may have accrued to the first of the steps. The general idea is that consecutive steps in a single plan eliminate uncertainty about the interpretation of the first step. Given these rules, the *combined* endorsements for the plan1 interpretation of the inputs **a**, **b** and the plan2 interpretation of the inputs **a**, **b**, **d** are derived from the endorsed step linkages shown above:

<u>plan1 interpretation of **a**, **b**:</u>	<u>plan2 interpretation of **a**, **b**, **d**:</u>
(**a** only grammatical possibility +)	(**b** other grammatical possibility -)
(**a** **b** continuity is desirable +)	(**d** only grammatical possibility +)
(**b** could be a mistake -)	(**b** **d** continuity is desirable +)
(**b** other grammatical possibility -)	(**d** could be a mistake -)

Note that (**a** could be a mistake -) has been erased by application of SCR1 for the plan1 interpretation, and that (**b** could be a mistake -) and (**a** **b** discontinuity is undesirable -) have been erased by SCR1 and SCR2 respectively for the plan2 interpretation.

4.3 Strengthening Endorsements

The semantic combining rules discussed above are unintuitive because they eliminate endorsements entirely, rather than increasing or decreasing the weight of endorsements (e.g., a more intuitive version of SCR1 should reduce the concern that a plan step is a mistake, not drop it entirely). Currently, we use numerical weights to reflect the strengths of endorsements, and adjust the weights to reflect combinations of endorsements. Since we are concerned that these numbers should mean the same under combination as combinations of endorsements, we have strictly limited ourselves to a single case of combination, namely *corroboration* of endorsements. We have identified three general situations where endorsements corroborate, that is, where two endorsements combine to create another "weightier" endorsement:

1. Corroboration of multiple instances of the same endorsement within a single plan step. For example, if an ambiguous plan step could continue one plan and start numerous others, then the weight of the "continuity is desirable" endorsement is greater than it would be if the step could continue a plan and start but a single plan.

2. Corroboration of instances of different endorsements of the same sign (both positive or negative) within the same plan step, resulting in a kind of synergetic increase in the belief in an interpretation. For example, the two negative endorsements "discontinuity is undesirable" and "other grammatical possibility" have a combined weight which is greater than the sum of their individual weights.

3. Corroboration of multiple instances of the same endorsement over consecutive plan steps. We believe in a plan more strongly if it is successively reinforced by the same positive endorsements. For example, we increase the weight of endorsements associated with a plan if the "continuity is desirable" endorsement appears in several consecutive steps.

4.4 Ranking Endorsements

We have said that the three components of semantics for endorsements are applicability conditions, combining rules, and ranking rules. We have explored two methods for ranking combinations of endorsements: one used the numerical weights of endorsements as described above, the other was a classification scheme to separate likely and unlikely alternatives.

We wanted combinations of endorsements to dictate at least a partial ordering on alternatives facing any decision-making program. We accomplished this in HMMM with a scheme for classifying step linkages as *likely*, *unlikely*, or *neutral*,[3] contingent on the presence of particular endorsements or combined endorsements. For example, a sufficient condition for being considered "likely" might be corroboration of two different, positive endorsements, and the condition for "unlikely" might be any negative endorsement. Interpretations can be ranked by assigning them to one of these implicitly-ordered classes, based on their endorsements. We think this kind of classification scheme can serve as a general model for ranking endorsements, since the criteria for membership in classes are flexible (and may be set dynamically); and since the number of classes is also flexible, ensuring adequate discrimination of alternatives. (The classification scheme was originally devised for a planning program which predicts a planner's next move to be from the class of "likely" moves.)

5.0 Discussion

The HMMM program raises many questions about endorsement-based reasoning. Two we did not address in the body of this paper concern the subjectivity and cost of endorsement-based reasoning.

Subjectivity of endorsements. Endorsement-based reasoning is not normative or prescriptive: there's no "correct" set of endorsements for a domain, no correct method for combining the endorsements of successive pieces of evidence. The endorsements discussed in this paper seem appropriate to the domain of plan recognition. We believe that ambiguity of plan steps reduces certainty in all interpretations of those steps, just as certainty is increased when two or more consecutive steps are interpreted as belonging to the same plan. Other people might design a different set based on their perceptions of the domain. The point is that this paper provides a framework for endorsement-based reasoning, but it is not prescriptive.

How much is required? The simple plan recognition example required few endorsements and only two semantic combining rules. We need more of each to handle other kinds of uncertainty and other relationships between endorsements. The number of endorsements and combining rules required for a domain depends on what you intend to do with them. If you wish to represent the *major* sources of uncertainty in a domain (e.g., the possibility of mistakes, ambiguity, disruption of an established scheme, etc.), then we believe the number of combining rules will be small. This is the approach we took for plan recognition. We expect that

[3] These terms are the names of classes; membership in any class is determined by endorsements. We imply no probabilistic interpretation of these terms.

endorsements can constitute a small investment for system-builders with a large payoff in terms of explanatory power and facilitation of knowledge engineering (since the expert can give reasons for uncertainty instead of numbers).

6.0 Conclusion

We propose endorsement-based reasoning as an explicit means of recording and combining reasons to believe and disbelieve propositions. The semantics of endorsements are specified by their applicability conditions, combining rules, and ranking rules. This paper addresses judgment -- assessment of support for propositions -- and thus the first two components of endorsement semantics. Decision, for which ranking of endorsements is needed, is a major concern of current work.

7.0 References

Buchanan, B. G., and Shortliffe, E. H., 1984. Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project. Chapter 10. Reading, Mass.: Addison-Wesley.

Carver, N. F., Lesser, V. R., and McCue, D. L., 1983. Focusing in plan recognition. Proceedings of the American Association for Artificial Intelligence.

Cohen, P. R., 1983. Heuristic reasoning about uncertainty: An artificial intelligence approach. TR-STAN-CS-83-986, Dept. of Computer Science, Stanford University. Also in press by Pitman Press.

Gadsden, J. A., 1984. An expert system for evaluating electronic warfare tasking plans for the royal navy. Proceedings of the First Conference on Artificial Intelligence Applications. IEEE Computer Society Press.

Tversky, A. and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. Science, 185, 1124-1131