From Proceedings of the 1996 Midwest AI and Cognitive Science Conference

Evidence Reliability in Nonmonotonic Domains and Nearly Monotonic Problems*

Norman Carver
Computer Science Department
Southern Illinois University
Carbondale, IL 62901
(carver@cs.siu.edu)

Abstract

Many AI domains are nonmonotonic. The resulting uncertainty is often dealt with by associating degree-of-belief ratings (conditional probabilities) with propositions. However, in some domains, a key issue is determining the subset of the available data that the agent should process to use as the basis for making a decision about how to act. Since belief ratings provide no sense of their "reliability" (if additional data were considered), they do not help an agent determine when sufficient data has been considered. We believe that in many domains, there is a strong correlation between current belief and ultimate belief once some fraction of the available data has been processed. We call such problems *nearly monotonic*, and suggest that they are important in that they can support efficient satisficing problem solving.

* This work was supported in part by the Department of the Navy, Office of the Chief of Naval Research, under contract N00014-95-1-1198. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

1 Introduction

Nonmonotonicity is a fact of life in many AI domains: what looks correct at one moment, may not as more information is gathered and assessed. Nonmonotonicity results in uncertainty about the state of the world and uncertainty about the actions that should be taken by an agent. One popular approach to representing and reasoning about this uncertainty is probability theory. Instead of just saying that some proposition (hypothesis) H is "likely," we assign a degree-of-belief to the proposition, which is the conditional probability of H given the current evidence E (i.e., $P(H \mid E)$). Knowing the degree-of-belief in its propositions can help an agent determine how it should act. For example, these values, along with a utility function, can be used in a decision-theoretic framework to select the optimal (rational) action to take.

However, in certain kinds of problems, so much data can be available to an agent that it is not practical to collect and process all of this data before making a decision. This is because the time it takes to select an action often has a negative effect on the utility of the action. Thus, agents must consider what data they will collect and process to use in making a decision.[1] The problem with conditional probability information in these situations is that it does not help an agent decide when enough data has been processed. For instance, just because the degree-of-belief in a proposition is 0.99 after some data has been processed, it may become 0 when the *very next piece* of data is processed—if the domain is truly nonmonotonic. In other words, if the complete set of data available to an agent is $\mathcal{D}$ and D is some proper subset of $\mathcal{D}$, then $P(h \mid D)$ alone tells us *nothing at all* about $P(h \mid \mathcal{D})$ (since $P(h \mid D)$ can be very high but $P(h \mid \mathcal{D})$ very low, or vice versa).[2]

What is needed is some measure of the "reliability" of the current belief rating. For example, how likely this rating is to be "close" to the *correct* value (the value obtained if the entire data set was processed). Furthermore, if such a measure is to be useful, nonmonotonicity cannot be "completely unpredictable." There must be a strong *correlation* between current and correct belief ratings, once the current ratings are based some reasonable fraction of the available data. In other words, ratings based on a portion of the data must be a reliable predictor of the correct ratings. We call problems with this general characteristic *nearly monotonic*. For such problems, detailed knowledge about the "reliability" of its current beliefs allows an agent to assess the quality of its solutions and limit the data it processes.

[1]More generally, there will be limits on all aspects of the reasoning that an agent might do in reaching a decision—i.e., rationality is *bounded/limited*. Here, we limit our discussion to the consideration of information that must be collected and/or then processed by the agent before it can be used in any reasoning—e.g., as in sensor interpretation for robot map making.

[2]For instance, if all you know is that your yard is wet, you may consider it very likely that it has recently rained. However, once you also know that your spouse just watered the lawn, the likelihood of it having rained becomes very low. What one thinks about rain knowing only about yard wetness does not predict what one will think about rain knowing about yard wetness and sprinkler usage.

Near monotonicity is particularly relevant to domains in which we will accept approximate, satisficing behavior rather than optimal, rational action.

The concept of a near monotonicity is explored further in the next section. This issue has arisen out of our research on distributed situation assessment and Section 3 contains a discussion of its role in these problems. The paper concludes with a summary of our conclusions and future plans.

2 Nearly Monotonic Problems

The basic idea behind the concept of nearly monotonic problems is that while a domain may be nominally nonmonotonic with increasing evidence, in nearly monotonic problems, once certain conditions have been achieved these problems behave as if they are nearly or approximately monotonic. For example, in vehicle monitoring, once a vehicle track hypothesis achieves a high degree-of-belief, it is *unlikely* that that this belief will be signficantly reduced by additional evidence and it is unlikely that the hypothesis will not be part of the correct interpretation.

It is impossible to give a single definition for near monotonicity. What we will instead do is present some possible characterizations of evidence "reliability" that we have been considering in our research on distributed situation assessment. These formulas are useful for assessing and using near monotonicity in that domain. The following notation will be used: $\mathcal{D}$ is the complete set of available data; D is some subset of $\mathcal{D}$; $BEL(H)$ is the current belief in proposition/hypothesis H given that data set D has been processed (it is $P(H \mid D)$); $BEL^*(H)$ is the "correct" belief in H given all of the data $\mathcal{D}$ (it is $P(H \mid \mathcal{D})$); MPE is the current MPE solution[3] given data set D; and MPE^* is the "true" MPE solution for data set $\mathcal{D}$.

The basic approach we take is to statistically characterize the properties that hypotheses or solutions would have considering the complete data set, given particular current characteristics. Five of the characterizations that we have considered are:

1. a conditional probability density function (conditional pdf) $f_{BEL^*|x}(b)$ that describes the probability that the ultimate belief in H is b given that the current belief is x, defined such that $\int_{p_1}^{p_2} f_{BEL^*|x}(b)db = P(p_1 \leq BEL^*(H) \leq p_2 \mid BEL(H) = x, ...)$;

2. the probability that the correct belief in the H will be greater than or equal to its current belief, $P(BEL^*(H) \geq BEL(H) \mid BEL(H) = x, ...)$;

3. the probability that the correct belief in H will be greater than or equal to some specified level l, $P(BEL^*(H) \geq l \mid BEL(H) = x, ...)$;

[3] MPE stands for *most probable explanation*—the most likely interpretation.

4. the probability that the hypothesis will ultimately be in the MPE, $P(H \in MPE^* \mid BEL(H) = x, ...)$;
5. the probability that the hypothesis will ultimately be in the MPE given that it is in the current MPE, $P(H \in MPE^* \mid BEL(H) = x, H \in MPE, ...)$.

A nearly monotonic problem would be one in which, once a proposition/hypothesis has certain characteristics (based on only a fraction of the available data), these probabilities will be *high enough* that it is appropriate to assume the proposition is "correct" without having to process additional data. The value of having such information is that an agent could assess the reliability of the data it is using in its decisions and process only the amount of data needed to achieve a sufficient level of reliability.[4] The above formulas assume that current degree-of-belief is an important factor in assessing when belief ratings are reliable predictors of correct belief, but they leave open what additional factors might be used. This will certainly depend on the particular problem domain, as the predictiveness of different characteristics is likely to vary across domains and systems will vary in their solution quality requirements. From our experience with situation assessment it appears that both current hypothesis belief and hypothesis type are important factors, but there are other factors that we are considering as well (e.g., some measure of the "quantity" of data supporting the hypothesis or the fraction of the relevant data that has been processed).

The concept of a nearly monotonic problem is related to some common notions in probability theory and decision theory. For example, belief networks can be used to do *hypothetical reasoning*: assessing what effect instantiating particular evidence nodes would have on the belief in proposition nodes. Clearly, this could provide the sort of reliability information we desire. The problem is that while this would work for small, finite networks, it would not work for domains like situation assessment involving large, indeterminate amounts of data/evidence. Likewise, the notion that problems become nearly monotonic in their behavior once a certain amount of data/evidence has been accumulated can be expressed in terms of *the value of information*. Basically, while additional data will typically improve an agent's decision (data has value), beyond some point additional data is unlikely to change the decision or at least make it "significantly better" (data has no/little value). When the negative effect of the time to reach a decision is factored in, we will reach a point where processing additional data has a negative utility. Again, though, it is rarely practical for an agent to use information value theory when deciding whether to process additional data or not. Thus, near monotonicity may not be an entirely novel concept, but there is still a need to develop practical mechanisms for exploiting it.

[4] Notice that we are talking about approximate problem solving here, since these correlations will generally be uncertain (< 1.0). Agents will sometimes erroneously assume propositions are correct when they are not or will reason with probabilities that are significantly incorrect.

3 Distributed Situation Assessment

We have been working to develop agent architectures and coordination strategies that result in distributed situation assessment (DSA) systems[5] with predictable performance characteristics (solution quality and communication loads) [2, 3, 4, 5]. In examining strategies that have been used in some past DSA systems, we discovered that "completely unpredictable" nonmonotonicity should lead to completely unpredictable solution quality results [3, 5]. The apparent success of systems using these strategies led us to look for an explanation. We believe that a key reason is that in many DSA domains, the problems become nearly monotonic once a portion of the data has been interpreted.

In DSA systems, we are interested in limiting the information that must be communicated among the agents. The ability to *efficiently* integrate (possibly interdependent) local agent solutions so as to end up with an appropriate global solution is a key issue for DSA. One strategy that some previous DSA systems [1, 6, 7] have used to achieve efficient operation is what we term the *consistent local solutions strategy* [4, 5]: agents transmit their local solutions, check the consistency of these solutions, and "merge" them when they are consistent. Because interpretation hypotheses are abstractions of the raw sensor data, if agents can communicate mainly solutions they can greatly reduce their communications. The advantage of having a nearly monotonic DSA problem is that consistency of local solutions can now be *highly predictive* that the merged, global solution will be correct, in the sense of being globally "best." Thus, when local solutions are consistent (and have appropriate characteristics) they can be cheaply merged while still ending up with a high quality global solution.

We have not yet shown that real-world SI problems are indeed nearly monotonic because of the difficulties of obtaining suitable data. As a first step toward this goal, [5] uses a recently developed framework for analyzing interpretation domans and problem solvers to provide support for the existence of nearly monotonic DSA problems. The *Interpretation Decision Problem* (IDP) formalism and system [8] has been used to analyze models of the Distributed Vehicle Monitoring Testbed (DVMT) [7], a DSA research system. Statistics gathered on these models clearly support the notion that SA problems can be nearly monotonic under reasonable conditions. For example, it was found that some types of interpretation hypotheses are very poor predictors of ultimate correctness (correct solution membership), regardless of their degree-of-belief. Certain other types of interpretation hypotheses, though,

[5] Situation assessment (SA) involves the determination of high-level, abstract *explanations* of sensor and related data. For example, in vehicle monitoring applications this involves tracking and identifying vehicles, and possibly determining the purpose of individual vehicles and patterns of vehicles. The *solution* to an SA problem is the "best" explanation for the data, according to some criteria. In a DSA system, each agent monitors only a portion of the overall "area" of interest, so agents must combine their solutions in order to construct a *global solution*.

were found to have a strong correlation between belief and correctness, and this occurred with relatively limited amounts of data. For instance, track hypotheses whose belief ratings were approximately 0.7, had a 0.92 probability of being correct.

4 Conclusion

In this paper, we have discussed the need for assessments of the reliability of degree-of-belief ratings for agents acting in nonmonotonic domains, which are faced with large amounts of data that they might make use of in choosing actions. We have also suggested that in some such domains there can be a strong correlation between the belief ratings calculated from a portion of the available data and the "correctness" of the propositions/hypotheses. Problems with this property were referred to as nearly monotonic.

We have found that this property is potentially important for DSA systems. It allows agents to use efficient (low communication cost) coordination strategies and still produce reasonable global solutions (as long as the local solutions meet certain criteria). This work has furthered our understanding of when distributed SA approaches are practical and what appropriate coordination strategies are. Near monotonicity does not guarantee efficiency, however, and distributed SA systems can sometimes be efficient even without the property. A focus of our future DSA research will be assessing the occurrence of near montonicity in real-world SA domains and determining its importance for efficient distributed SA.

One general lesson from this work is that while nonmonotonicity is a fact of life in many AI domains, it is often not completely arbitrary or unpredictable, and models of its characteristics might yield important benefits. If problem solving must be approximate and satisficing, a statistical characterization of the montonicity/nonmonotonicity in the domain makes it possible to evaluate the reliability of approximate solutions based on incomplete data. The work can also be viewed as part of a general trend in AI of trying to relate system performance and domain characteristics to better understand what architectures are appropriate for what domains. A particular approach does not just *happen* to work well/poorly in some domain. There are reasons why this is the case, which can be understood by characterizing the properties or structure of different domains.

References

[1] Norman Carver, Zarko Cvetanovic, and Victor Lesser, "Sophisticated Cooperation in FA/C Distributed Problem Solving Systems," *Proceedings of AAAI-91*, 191–198, 1991.

[2] Norman Carver and Victor Lesser, "A First Step Toward the Formal Analysis of Solution Quality in FA/C Distributed Interpretation Systems," *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence*, July, 1994.

[3] Norman Carver, "Examining Some Assumptions of the FA/C Distributed Problem-Solving Paradigm," *Proceedings of the Midwest Artificial Intelligence and Cognitive Science Society Conference*, 37–42, April, 1995.

[4] Norman Carver and Victor Lesser, *A Formal Analysis of Solution Quality in FA/C Distributed Sensor Interpretation Systems*, Technical Report 95-05, Computer Science Department, Southern Illinois University, 1995 (submitted for publication).

[5] Norman Carver, Victor Lesser and Robert Whitehair, *Nearly Monotonic Problems: A Key to Effective FA/C Distributed Sensor Interpretation*, to appear in *Proceedings of AAAI-96*, 1996.

[6] Victor Lesser and Daniel Corkill, "Functionally Accurate, Cooperative Distributed Systems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 1, 81–96, 1981.

[7] Victor Lesser and Daniel Corkill, "The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem Solving Networks," *AI Magazine*, vol. 4, no. 3, 15–33, 1983 (also in *Blackboard Systems*, Robert Engelmore and Tony Morgan, editors, Addison-Wesley, 1988).

[8] Robert Whitehair, *A Framework for the Analysis of Sophisticated Control*, Ph.D. Thesis, Computer Science Department, University of Massachusetts, 1996.