Approximate Sensor Interpretation

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

Abstract

Sensor interpretation (SI) involves determining abstract explanations for sensor data. SI differs in several significant ways from the kind of "diagnosis problems" that have been heavily studied within the belief network community. These differences lead to the need for approximate, satisficing problem-solving techniques in most real-world SI problems. Currently, there are no AI techniques with well understood properties that can apply a wide range of approximate SI strategies. In this paper we will examine the differences between SI and diagnosis that lead to the need for approximation, and discuss several approximation techniques. We will then consider the two main AI approaches to SI, blackboard systems and dynamic belief networks, and explore their deficiencies for SI. As a point of comparison, we will also consider techniques used by the target tracking community.

Introduction

Sensor interpretation (SI) involves the determination of abstract, conceptual *explanations* of sensor data and/or other information.[1] In other words, a description of the "events" in the environment that are responsible for producing the sensor data. A wide range of tasks can be viewed as SI problems: vehicle monitoring and tracking, robot map making, sound understanding for robotic hearing (auditory scene analysis), speech understanding, and so forth.

SI differs in some important ways from the types of problems that have been highly studied by the belief network community. For example, in SI problems, there are typically an indeterminate number of events that may be occuring in the environment and systems are faced with *data association uncertainty*. These characteristics require the use of constructive problem solving techniques and the potential for exponential growth in the size of the "belief network" as data is processed. As a result, approximation is required for most real-world SI problems. For example, the *target tracking* literature contains a variety of approximation approaches that are used in conjunction with the basic Kalman filtering techniques.

In the AI community, *blackboard systems* have been a popular framework for tackling complex SI problems. A key reason for this is that the blackboard model supports a variety of SI approximation techniques. In particular, blackboard-based SI systems typically employ a sophisticated search through partial interpretation states. What blackboard-based SI sacrifices is the ability to make formal statements about the quality of its solutions and the characteristics of its control decisions. The belief network community has so far focused on *dynamic belief networks* (DBNs) as the main approach to SI problems. One way DBNs differ from most blackboard systems is their ability to perform exact probabilistic inference and select optimal interpretations. Since they do not support key SI approximation techniques, however, they are unlikely to be useful for all but very simple SI problems.

Clearly, what we would like is to have the best of both approaches: the flexibility and approximation techniques of blackboard systems, with the ability to make formal statements about solution quality and the properties of the control decisions. Unfortunately, there are not currently AI approaches to SI with all of these characteristics. We are currently exploring how to formally characterize the properties of interpretations developed by the (blackboard-based) RESUN SI framework. We are interested in adapting ideas from the belief network community about approximate probablistic inference. A critical issue is that "belief nets" for SI problems will typically need to be incomplete due to the potential for exponential network growth. We are considering models of the properties of SI domains to deal with incomplete networks.

In this paper, we will first examine the characteristics of SI that make it a difficult problem and talk about ap-

[1] More generally, we may speak of *situation assessment*, especially if we include non-sensor data/evidence.

proximate strategies for SI. As a point of comparison, we will then look at the sorts of techniques that are used by the target tracking community. This is followed by introductions to the use of blackboard systems and dynamic belief networks for SI—their strengths and weaknesses. The paper concludes with a brief summary and a short discussion of our current research interests.

Interpretation vs. Diagnosis

Interpretation problems differ in several significant ways from the kinds of problems that have typically been studied by the abductive inference and belief network communities (e.g., see (Peng and Reggia 1990) and (Pearl 1988), respectively). For simplicity, we will refer to these as *diagnosis problems*. Because of the differences, techniques that are appropriate for diagnosis problems may not be sufficient for SI problems. In this section, we will first define some terminology and then examine the differences between interpretation and diagnosis.

As we have said, the point of interpretation is determining an *explanation* for sensor data and other information. An *interpretation* is a set of *hypotheses* about "events" that might have caused the data (and so explain it). For example, events could be vehicles moving through the environment or the speaking of words/sentences. Each hypothesis explains some subset of the data and together an interpretation's hypotheses explain all of the data. Typically, we are interested in interpretations whose hypotheses are from a subset of the abstraction types—what Pearl has termed the *explanation corpus* (Pearl 1988). The process of interpretation is based on a causal model that relates data characteristics to types of "events." An interpretation system uses this model to make *abductive inferences* that identify possible explanations (causes) for the data (Carver and Lesser 1991).

Interpretation is an inherently uncertain process. In general, there will be multiple possible interpretations of any data set—i.e., multiple alternative sets of "events" that could have caused the data. In addition, many interpretation domains involve significant "noise" in the data from things like sensor errors, environmental factors, and so forth. Because of this uncertainty, there must be some way to assess the strength of the evidence for the alternative interpretations. The *solution* to an interpretation problem is the interpretation that is judged "best" according to some criteria. In a probabilistic context, one possible definition of best is the *most probable explanation* (MPE) (Pearl 1988).[2]

So far, diagnosis problems sound very similar to SI problems: they involve determining explanations (e.g., diagnoses) for evidence/data, diagnosis is based on a causal model, and there can be multiple, alternative explanations that must be considered. The key characteristics of diagnosis problems, though, are that they have a *fixed* set of interpretation hypotheses (e.g., diseases) and data nodes, with *known, fixed* relations among them. In other words, a complete and static (probabilistic) model is available. Probabilistic inference in networks for such problems has been studied extensively. The conditional probability of nodes as well as the MPE can be determined by plugging the available data into evidence nodes and appropriately propagating its effects (Pearl 1988; Russell and Norvig 1995). While inference has been shown to be NP-hard for general network topologies (Cooper 1990; Shimony 1994), efficient exact and approximate techniques have been developed that can handle many reasonable size problems.

The primary way that interpretation problems differ is that they lack complete and static models that connect the data with possible explanations. While interpretation problems have a fixed set of interpretation and data *types*, they can have an *indeterminate number* of *instances* of any of these types.[3] For example, in a vehicle monitoring system, an unknown number of vehicles will have been responsible for the overall data set and each vehicle will produce a "track" of an unknown number of sensor data points. The causal model used in the interpretation process identifies the connections between data types and explanation types, but not between *instances* of these types.

As a result, the associations between the individual pieces of data and individual interpretation hypotheses are a priori *unknown*. This leads to what is known in the *target tracking* literature as the *data association problem* (DAP) (Bar-Shalom and Fortmann 1988): which target should data be associated with? The DAP gives rise to what has been termed *correlation ambiguity* or *origin uncertainty*: it is ambiguous/uncertain which potential hypothesis each piece of data should be associated with (and provide evidence for). Diagnosis problems simply do not involve the DAP or correlation ambiguity.

The DAP and the possibility of an indeterminate number of event instances, can lead to a combinatorial explosion in the number of possible interpretations for a data set. For example, in a vehicle monitoring system, every single piece of data potentially could have come from: (1) any already hypothesized vehicle, (2) noise/clutter, or (3) a new (previously undetected) vehicle. Unless it is possible to *conclusively* rule out many of these interpretations, the number of hypotheses will

[2] The MPE is also called the *maximum a posteriori probability* (MAP) interpretation.

[3] By indeterminate, we mean both that the number is a priori unknown and that it changes over time.

grow exponentially with the amount of data examined.

To help understand the implications of the DAP, consider a medical diagnosis problem involving data association uncertainty: A doctor has a set of patients and a set of test results. However, the tests have not been labeled by patient and some may even be for patients that he has not seen yet (and knows nothing about). What is the diagnosis for each of the tested patients? This is clearly a much more difficult problem than conventional diagnosis problems since the doctor not only must diagnose the patients, he must figure out how many patients tests he has and associate tests with patients.[4]

Interpretation systems also face problems not faced by most diagnosis systems that arise from the nature of their sensor data evidence. First, data from the same sensor over time may not be conditionally independent (given an interpretation hypothesis). Whether this is the case or not will depend on how much detail we represent in the interpretation hypotheses and the level of detail in the causal model. The second problem with sensor data is that there can be a massive amount of it when there are multiple passive sensors, continuously operating in a noisy environment; too much to completely process. On the other hand, in many interpretation problems we do not need to have explanations of every piece of data since we only care about certain events or interpretation types (e.g., platforms/targets vs. clutter). Finally, SI data and hypotheses typically involve multiple, often *continuous-valued* attributes.

Yet another source of difficulty for many SI problems is the need for real-time performance. This is true, for example, in many target tracking applications, where systems must constantly provide current interpretations. An interesting issue is that R/T constraints may be different for different components of the overall interpretation since they can be a function of the type of "event." For instance, it is often more critical to rapidly confirm the presence and state of hostile platforms than of friendlies.

Approximate Interpretation

To deal with the characteristics discussed in the previous section, interpretation systems must be *constructive* (Carver and Lesser 1991; Clancey 1985) and they also must usually make use of *approximate, satisficing strategies* to determine solutions. There are seven basic approximation techniques that can be used by interpretation systems:

1. process only part of the available data;
2. construct only some of the possible interpretation hypotheses for the processed data;
3. periodically delete (prune) unlikely interpretation hypotheses;
4. compute approximate belief ratings (conditional probabilities) for the hypotheses (perform limited "evidence propagation");[5]
5. compute beliefs only for certain types of interpretation hypotheses (e.g., those from the explanation corpus);
6. consider only some of the possible interpretations (hypothesis combinations);
7. use criteria other than the MPE to select the solution (e.g., assemble solutions from hypotheses whose belief ratings surpass some *acceptance threshold*).

Obviously, these techniques are not independent. If a system does not process all of the data then it cannot in general create all possible interpretations of the complete data set nor compute the true conditional probabilities of the hypotheses. If a system does not create every possible interpretation hypothesis for some data, this not only limits the interpretations that can be considered, it also results in hypothesis belief ratings being only approximations of the true conditional probabilities of the hypotheses. This is because incomplete hypothesis construction results in incomplete propagation of the effects of evidence. Figure 1 provides an example of this situation. The bottom line is that these approaches will result in interpretation solutions that are only approximations of the optimal, MPE interpretation: they may be incomplete or they may not be the most probable composite interpretation.

Despite the obvious drawbacks of approximate approaches, many real-world SI problems simply require approximation. Furthermore, there are reasons why approxmiate SI can often be effective: systems are typically interested in only certain types of phenomena out of all the environmental phenomena for which there are models (e.g., targets/platforms vs. noise/clutter); data may be redundant due to the existence of multiple sensors; it may not be necessary to process all relevant

[4] To be fair, this example somewhat overstates the difficulties typically caused by the DAP. In problems like vehicle monitoring, basic consideration of the laws of physics can eliminate (or make extremely unlikely) many possible associations simply by virtue of position information, but this is less so for medical diagnosis (though some diseases are highly correlated to basic patient factors like age). Another reason is that in many interpretation problems, it is reasonable to assume that a single source was responsible for each piece of data—unlike medical diagnosis where the simultaneous occurrence of multiple diseases with overlapping symptoms is a key issue.

[5] We are using "evidence propagation" in the same basic sense that (Pearl 1988) refers to "belief propagation." We will use the term "belief" to mean the degree-of-belief accorded to a hypothesis or interpretation. This will be either the conditional probability or approximate conditional probability of the object.

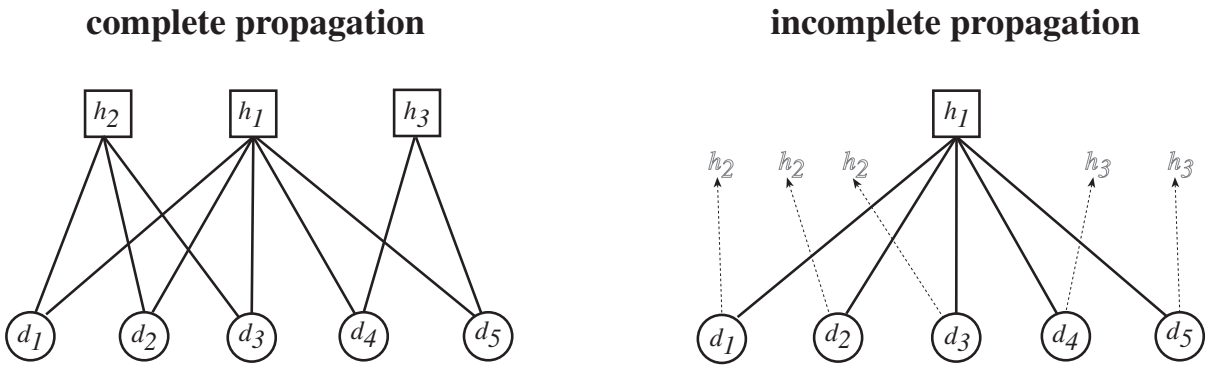

Figure 1: Approximate interpretation due to incomplete hypothesis construction and propagation.
In the complete construction/propagation case, all possible (alternative) interpretations h_1, h_2, and h_3 have been created (using the most complete support possible). This allows the system to determine the conditional probability of h_1 given the available data $\{d_1, \ldots, d_5\}$. In the incomplete construction/propagation case, the alternative explanations to h_1 have not been created. This means that the belief computed for h_1 can be only an approximation of the true conditional probability since the likelihood of the alternative explanations can be only approximately considered.

data and make every evidential inference in order to be sufficiently certain of interpretations; and so forth.

A key issue for approximate SI is determining appropriate approximation strategies. Ideally, this should be done dynamically, in response to changing interpretations, data loads, ambiguity, and goals. Many approximate strategies are incremental in nature, such that additional time can be used to improve the interpretation. For example, if we are constructing only the "more likely" interpretation hypotheses, additional time can be used to construct progressively more hypotheses and so allow for improved belief computations. Likewise, additional time can be used to process increasingly larger subsets of the available data. Thus, many of these strategies are consistent with the notion of *flexible computations* or *anytime algorithms* (Dean and Wellman 1991).

Target Tracking

Target tracking is a much studied SI application. It involves identifying "targets" using "measurements" derived from sensor data, and correlating measurements over time to track the targets (and perhaps maintain other state information). Target tracking is a specialized SI application because limited interpretation of the data is being done—e.g., there is typically no attempt to explain target behavior or the patterns of multiple targets. The need for approximate techniques has been well known within this community to deal with data association uncertainty, large numbers of targets, high data loads, maneuvering targets, and so forth. An excellent introduction to the issues and standard techniques in target tracking can be found in (Bar-Shalom and Fortmann 1988).

We will consider the Bayesian Multiple Hypothesis Approach (BMH) of (Cox and Leonard 1994), since it is one of the more advanced and flexible target tracking algorithms. On each sensor cycle, BMH first forms all possible combinations of the new data with: (1) each of its existing track hypotheses, (2) the false alarm/clutter interpretation, and (3) the new track interpretation. Once this is done, exact conditional probabilities are computed for each of the possible associations (hypotheses), using a Kalman filter approach (with certain assumptions about the distributions of noise and new features). BMH deals with the DAP by considering all possible associations between the data and possible hypotheses on each cycle and computing conditional probabilities. To deal with the exponential growth in the number of hypotheses that would be generated over time, DMH resorts to approximation. Periodically, low likelihood hypotheses are *pruned*, so that they are not used in forming further associations.

While pruning can be based on conditional probabilities given the previous model, each pruning effectively results in the state that will be used in the next cycle of the Kalman filter being only an approximation of the true state. The formal properties of this approximation are not discussed in (Cox and Leonard 1994) and are apparently unknown. One reason this may not be considered a serious problem is that it appears common to assume that uncertainty can be fairly rapidly resolved—i.e., there will be little uncertainty about a target after the fusion of data from a relatively small number of time slices. In fact, many target tracking approaches do not even maintain multiple hypotheses between cycles, they simply select the most likely on each cycle.

Another limitation of BMH is that it has little flexibility: it works time slice by time slice, evaluating all

data and creating all possible interpretations (given its previous model). To allow this type of processing to occur in real-time requires that the number of targets and amount of data per time slice not become excessive, and that the number of hypotheses being maintained not be allowed to grow too large.

Blackboard-Based SI

The blackboard model of problem solving grew out of the Hearsay-II (HSII) speech understanding system (Erman et al. 1980) and was designed to deal with the difficult problems of SI. Among the key ideas behind this model are that problem solving should be both *incremental* and *opportunistic*. That is, solutions should be constructed piece by piece and at different levels of abstraction, working where the available data and intermediate state of problem solving suggest the most progress can be made. More substantial introductions to the blackboard model can be found in (Carver and Lesser 1994; Engelmore and Morgan 1988) and in (Carver and Lesser 1992), which concentrates on blackboard-based SI. A recent blackboard-based SI system is described in (Lesser et al. 1993).

In the basic HSII model, a blackboard system is composed of three main components: the *blackboard*, a set of *knowledge sources* (*KSs*), and a control mechanism. The blackboard is a global database (i.e., shared by all the KSs) that contains the (sensor) data and interpretation hypotheses. The KSs embody the problem solving knowledge of the system: they examine the blackboard and can add, modify, or even delete hypotheses when appropriate. For SI problems, there would be KSs to identify possible explanations for data/hypotheses and construct hypotheses representing these explanations. The blackboard is typically structured as a set of *levels*. For interpretation problems, the blackboard levels are basically organized as a partial order, with data levels at the "bottom" and abstract explanation levels at the "top." Levels are themselves structured in terms of a set of *dimensions*, which are used to define the "location" of a hypothesis within a level. This makes it possible to provide efficient *associative retrieval* of hypotheses. The control mechanism decides what actions (knowledge source instantiations) the system should take next.

The blackboard approach to SI has emphasized the need for approximation and sophisticated searches to solve complex SI problems. This has produced systems with great flexibility. Blackboard-based systems are capable of implementing all of the approximation techniques listed earlier. Interpretation hypotheses can be constructed and refined incrementally, as part of a search process driven by "sophisticated control architectures." A variety of approximate knowledge sources can be defined and applied. Blackboard systems do not have to work time slice by time slice, forward in time; they do not have to create all interpretations of the data they process; and they do not have to process all the available data.

For example, blackboard systems can examine the data abstractly, looking for likely "targets," and then *selectively* processing data over a range of times to confirm/deny and refine their hypotheses (e.g., (Durfee and Lesser 1988)). Likewise, they can focus their activities on pursuing only interpretation hypotheses of most value—as in possible hostile aircraft vs. friendly aircraft. Because they incrementally develop hypotheses, blackboard systems can also work at multiple levels of abstraction—they need not immediately explain all data in terms of the ultimate (explanation corpus) interpretation types. This is one mechanism for dealing with the combinatorics of the DAP: implicitly representing uncertainty about higher level associations by not creating links representing those associations until sufficient data is acquired.

The blackboard model has also emphasized the need to dynamically and opportunistically adjust strategies. This means that blackboard systems can adapt their problem-solving strategies as data loads, ambiguity, system goals, and available time change. Decisions about what hypotheses should be pursued and how they should be pursued are based on the intermediate state of problem solving (the current hypotheses, data, goals, and so forth). This allows blackboard-based interpretation systems to deal with resource limitations by reasoning about appropriate solution quality vs. processing time trade-offs.

The limitations of blackboard-based SI systems have been and continue to be their lack of formalism. Blackboard systems have almost invariably used ad-hoc belief representations and solution selection strategies. This has made it impossible to determine the properties of the solutions produced by these systems (e.g., how they compare to the MPE solutions) other than through empirical means. Likewise, while blackboard systems can engage in sophisticated reasoning in making control decisions these approaches have not been formalized (e.g., in a decision-theorectic framework).

DBN-based SI

As we have shown, interpretation forces a system to deal with issues that are not raised by diagnosis problems. For example, since the number of possible data and interpretation instances are a priori unknown, interpretation problem solving is necessarily constructive. With a belief net, this means growing the network as data arrives and is processed. As each piece of data is processed, a corresponding evidence node could be created (with appropriate conditional probablity informa-

tion), new explanation nodes may need to be created, and evidential links added to connect the evidence node to nodes the data directly supports.

Instead of this approach, however, the belief net community has focused on *dynamic belief nets* (DBNs) to deal with the temporal issues raised by interpretation. The basic idea behind DBNs is to construct new instances of the dynamically changing portions of the belief net for each *time slice*,[6] but make use of the *Markov Property* to eliminate all but the latest two time slices of information by doing a *rollup* of all previous sensor information. Instead of having, say, *a single* vehicle hypothesis with supporting data/evidence nodes being added over time, a DBN would have a time t vehicle hypothesis, a time $t+1$ vehicle hypothesis, and so on. Each such vehicle hypothesis would be supported by data from its time slice plus by the previous *vehicle hypothesis*—not (directly) from all the accumulated data. Introductions to DBNs can be found in (Dean and Wellman 1991; Nicholson and Brady 1994; Russell and Norvig 1995).

The DBN approach to interpretation is quite *in*flexible. A DBN works time slice by time slice, doing complete and exact interpretation—i.e., determining all possible interpretations of the new data and computing exact probabilities. While a DBN can apply certain approximate inference techniques, there is no ability to selectively and opportunistically search. Russell and Norvig (Russell and Norvig 1995) (p. 518) state that "probably the most important defect of DDNs [dynamic decision networks] is that they retain the property of forward search through concrete [i.e., complete] states...." If we compare the DBN approach to the blackboard approach, we see that this is a key way that blackboards get their flexibility and power: they are not limited to forward search and can deal with partial states rather than complete states.

Another thing to note about DBNs is that they do not completely address the DAP and the resulting possibility of exponential growth in the number of interpretations and thus network structure. DBN's do reduce growth in the network that would come from adding numerous data/evidence nodes. *Rollup nodes* (Russell and Norvig 1995) are used to compactly represent previous sensor data, so that only the last time slice of the network needs to be maintained. This does not address the potential for exponential growth in the number of interpretation hypotheses over time, however. In general, approximation techniques will be required to deal with this issue. Furthermore, while the rollup approach works for simple interpretation problems (e.g., (Huang et al. 1994; Nicholson and Brady 1994)), it may not be practical for interpretation problems involving sensor data with multiple, continuous-valued attributes. For example, suppose that environmental characteristics in some region are causing particular distortions in the signal received from a target vehicle (e.g., certain frequencies shifted by x Hz.). Expectations about the continued appearance of this distortion need to be maintained between time slices. Doing this could require extremely complex, infinite-valued rollup nodes, which are not going to be practical in belief network computations.

[6] Each discrete time when new sensor data arrives.

Conclusion

The inherent properties of SI mean that resource limitations will be a critical issue for most real-world applications. Thus, approximate and resource-bounded problem solving techniques are required. Currently, there are no AI techniques for SI that can both provide these capabilities and have well understood properties. Blackboard systems can support approximate SI, but produce solutions whose quality cannot be formally assessed. They can dynamically/opportunistically consider how to proceed, but do not include a formal model for reasoning about resource trade-offs. DBNs can produce optimal, MPE solutions when applicable, but do not support most SI approximation techniques. Of course, this situation is not unique to AI approaches to SI. While target tracking algorithms are often based on formal techniques like the Kalman filter, they typically resort to approximations outside of the formal system (and they cannot reason about resource vs. solution quality trade-offs).

We are currently exploring how to formally characterize the properties of interpretations developed by the (blackboard-based) RESUN SI framework (Carver and Lesser 1991). Our intention is to try to adapt approximate probablistic inference techniques from the belief network community. An important issue here is that the research on approximate inference in belief nets has largely assumed that a complete network is available and the critical factor is that exact inference is not tractable. As we have discussed here, this assumption is often going to be invalid for SI because of the potential for exponential growth of the network. Approximation techniques for SI must be able to work with incomplete networks (i.e., partial states). It appears to us that the ability to make formal statements about solution quality in such situations will require knowledge of the characteristics of the domain. For example, we are considering properties such as *near monotonicity* (Carver, Lesser, and Whitehair 1996) to deal with evaluation in incomplete networks.

Another very exciting development that we intend to make use of is Whitehair and Lesser's *Interpretation*

Decision Problem (IDP) formalism (Whitehair 1996). This framework uses context free attribute grammars to represent complex SI domains and SI problem solvers. It allows one to compare the *inherent complexity* of interpretation in domains with particular characteristics and to analyze the appropriateness of a problem solving system for a domain. Unlike much formal work, it is powerful enough to represent real-world SI problems. So far, it has been used to analyze the properties of several sophsticated control strategies used in previous blackboard-based SI systems. It hold the promise of providing a method for formally assessing the characteristics of approximate, blackboard-based SI strategies. Unfortunately, the current IDP framework does not assess performance in terms of standard metrics like conditional probabilities or MPE solutions.

Acknowledgements

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.

References

Yaakov Bar-Shalom and Thomas Fortmann, *Tracking and Data Association*, Academic Press, 1988.

Norman Carver and Victor Lesser, "A New Framework for Sensor Interpretation: Planning to Resolve Sources of Uncertainty," *Proceedings of AAAI-91*, 724–731, 1991.[1]

Norman Carver and Victor Lesser, "Blackboard Systems for Knowledge-Based Signal Understanding," in *Symbolic and Knowledge-Based Signal Processing*, Alan Oppenheim and Hamid Nawab, editors, Prentice Hall, 205–250, 1992.

Norman Carver and Victor Lesser, "The Evolution of Blackboard Control," *Expert Systems with Applications*, special issue on The Blackboard Paradigm and Its Applications, vol 7, no. 1, 1–30, 1994.[1]

Norman Carver, Victor Lesser, and Robert Whitehair, "Nearly Monotonic Problems: A Key to Effective FA/C Distributed Sensor Interpretation?," *Proceedings of AAAI-96*, 88–95, 1996.[1]

William Clancey, "Heuristic Classification," *Artificial Intelligence*, vol. 27, 289–350, 1985.

G.F. Cooper, "The Computational Complexity of Probabilistic Inference using Bayesian Belief Networks," *Artificial Intelligence*, vol. 42, 393–405, 1990.

[1]Can be obtained from:
"http://www/cs.siu.edu/∼carver/".

Ingemar Cox and John Leonard, "Modeling a Dynamic Environment using a Bayesian Multiple Hypothesis Approach," *Artificial Intelligence*, vol. 66, 311–344, 1994.

Thomas Dean and Michael Wellman, *Planning and Control*, Morgan Kaufmann, 1991.

Edmund Durfee and Victor Lesser, "Incremental Planning to Control a Time-Constrained, Blackboard-Based Problem Solver," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 24, no. 5, 647–662, 1988.

Robert Engelmore and Tony Morgan, editors, *Blackboard Systems*, Addison-Wesley, 1988.

Lee Erman, Frederick Hayes-Roth, Victor Lesser, and D. Raj Reddy, "The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty," *Computing Surveys*, vol. 12, no. 2, 213–253, 1980 (also in *Blackboard Systems*, Robert Engelmore and Tony Morgan, editors, Addison-Wesley, 31–86, 1988).

T. Huang, D. Koller, J. Malik, G. Ogasawara, B. Rao, S. Russell, and J. Weber, "Automatic Symbolic Traffic Scene Analysis Using Belief Networks," *Proceedings of AAAI-94*, 1994.

Victor Lesser, Hamid Nawab, Izaskun Gallastegi, and Frank Klassner, "IPUS: An Architecture for Integrated Signal Processing and Signal Interpretation in Complex Environments," *Proceedings of AAAI-93*, 249–255, 1993.

Ann Nicholson and Michael Brady, "Dynamic Belief Networks for Discrete Monitoring," *IEEE Transactions on Systems, Man, and Cybernetics*, special issue on Knowledge-Based Construction of Probabilistic and Decision Models, vol. 24, no. 11, 1593–1610, 1994.

Judea Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.

Yun Peng and James Reggia, *Abductive Inference Models for Diagnostic Problem-Solving*, Springer-Verlag, 1990.

Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, 1995.

Solomon Shimony, "Finding MAPs for Belief Networks is NP-Hard," *Artificial Intelligence*, vol. 68, 399–410, 1994.

Whitehair, R. 1996. A Framework for the Analysis of Sophisticated Control. Ph.D. diss., Computer Science Department, University of Massachusetts.