A New Framework for Sensor Interpretation: Planning to Resolve Sources of Uncertainty

Norman Carver and Victor Lesser
Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01003
(carver@cs.umass.edu)

Abstract

Sensor interpretation involves the determination of high-level explanations of sensor data. Blackboard-based interpretation systems have usually been limited to incremental hypothesize and test strategies for resolving uncertainty. We have developed a new interpretation framework that supports the use of more sophisticated strategies like differential diagnosis. The RESUN framework has two key components: an evidential representation that includes explicit, symbolic encodings of the *sources of uncertainty* (*SOUs*) in the evidence for hypotheses and a script-based, incremental control planner. Interpretation is viewed as an incremental process of gathering evidence to resolve particular sources of uncertainty. Control plans invoke actions that examine the symbolic SOUs associated with hypotheses and use the resulting information to post goals to resolve uncertainty. These goals direct the system to expand methods appropriate for resolving the current sources of uncertainty in the hypotheses. The planner's *refocusing mechanism* makes it possible to postpone focusing decisions when there is insufficient information to make decisions and provides opportunistic control capabilities. The RESUN framework has been implemented and experimentally verified using a simulated aircraft monitoring application.

Introduction

Sensor interpretation involves the determination of high-level *explanations* of sensor and other observational data. The interpretation process is based on a hierarchy of abstraction types like the one in Figure 1 for a vehicle monitoring application. An interpretation system incrementally creates or extends *hypotheses* that represent possible explanations for subsets of the data. In vehicle monitoring, data from sensors (e.g., Acoustic Data and Radar Data) is abstracted and correlated to identify potential vehicle positions (Vehicle hypotheses), vehicle movements (Track hypotheses), and vehicle goals (Mission hypotheses). Interpretation can be difficult because there may be combinatorial numbers of alternative possible explanations of the data, creating each hypothesis may be computationally expensive, the correctness of the hypotheses will be uncertain due to uncertainty in the data and problem solving knowledge, and the volume of data may be too great for complete examination.

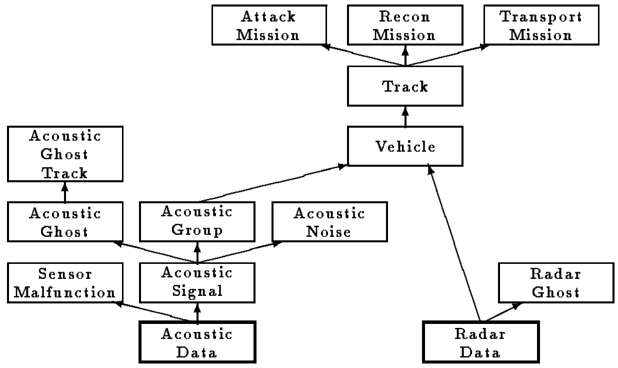

Figure 1: Vehicle monitoring abstraction hierarchy.

In order to understand the complexities of the interpretation process, it is useful to understand the distinction [Clancey 1985] draws between *classification problem solving* and *constructive problem solving*. In classification problem solving, the solution is selected from among a *pre-enumerated* set of all the possible solutions. In constructive problem solving, the set of possible solutions is determined as part of the problem solving process. While problems like simple diagnosis [Peng & Reggia 1986] can be approached using classification techniques, interpretation problems require constructive problem solving because of the combinatorics of their answer spaces. For example, in vehicle monitoring, an (effectively) infinite number of different Track hypotheses is possible, an indeterminate number of instances of Track hypotheses may be correct (since the number of vehicles that may be monitored is unknown), and correlation ambiguity produces a combinatorial number of data combinations to be considered. Clancey notes that constructive problem solving requires capabilities not required for classifica-

*This work was supported by the Office of Naval Research under University Research Initiative grant number N00014-86-K-0764.

tion problem solving—e.g., the ability to apply significant amounts of knowledge to focus the construction process. In addition, well-developed evidential reasoning techniques like Dempster-Shafer and Bayesian networks [Pearl 1988] are directly applicable only to classification problem solving [Carver 1990].

Interpretation problems have often been approached using *blackboard* frameworks. This is because the blackboard model supports constructive problem solving and because it supports opportunistic control for dealing with uncertain data and problem solving knowledge [Carver 1990, Lesser & Erman 1977]. Despite the power of the blackboard model, most blackboard-based interpretation systems (e.g., [Durfee & Lesser 1986, Erman et al. 1988, Lesser & Corkill 1983, Nii et al. 1982, and Williams 1988]) have been limited to using variations of *incremental hypothesize and test*[1] strategies for resolving interpretation uncertainty. The designers of the Hearsay-II architecture believed that blackboard systems would have the capability to do *differential diagnosis* because of their integrated representation of alternative, competing hypotheses [Lesser & Erman 1977]. However, explicit differential diagnosis techniques have not been exploited by blackboard-based interpretation systems[2] because of limitations in their evidence representations and control frameworks [Lesser & Erman 1977, Carver & Lesser 1990].

To illustrate the kind of control reasoning that an interpretation should be able to do, consider the following vehicle monitoring system scenario: "In order to meet its goals, the system has to reduce its uncertainty in hypothesis $Track_1$. To do this, it must first determine the reasons $Track_1$ is uncertain. Examining $Track_1$, it finds that a major source of uncertainty is the existence of a competing hypothesis, $Track_2$, which can provide an alternative explanation for a part of $Track_1$'s supporting data. In examining the reasons why $Track_2$ is uncertain, the system finds that the portion of $Track_2$'s supporting data which is not also supporting $Track_1$, might actually be able to be explained as a ghost. If this were the case, it would decrease the belief in $Track_2$ thereby helping to resolve the uncertainty in $Track_1$. For this reason, the system decides to construct a hypothesis representing the alternative ghosting explanation and then attempts to sufficiently prove or disprove it...."

This example shows that in interpretation problems, the process of making control decisions may require that a system be able to: gather information for the control process, consider the evidential relationships among hypotheses, understand how different methods can be used to satisfy its goals, and carry out methods for resolving uncertainty that involve sequences of actions. Sophisticated interpretation also requires the ability to do detailed control reasoning so that the actions taken can be very sensitive to the goals of the interpretation process (termination confidence criteria, time available, etc.) and the particulars of the situation (current evidence and uncertainties, data characteristics, availability of sensors, etc.). For instance, the failure to find data to extend a Track hypothesis might be due to the data having been missed by the sensor (due to some environmental disturbance, etc.). However, indiscriminately pursuing this possibility can lead to a combinatorial explosion in the number of hypotheses being considered. Thus, the decision about how to resolve the Track's uncertainty must carefully consider whether it is currently appropriate to pursue the possibility of missing data; even if it is because this possibility that the Track continues to be pursued, it may be appropriate to look at alternative sources of uncertainty first.

In this paper, we will describe a new interpretation framework called RESUN. This framework supports the use of sophisticated interpretation strategies. It provides an alternative to conventional blackboard systems for interpretation. The RESUN framework is introduced in the next section. Its evidential representation system and control planner are described in more detail in the following two sections. The final section of the paper discusses the status of our research, presents some experimental results, and concludes with a summary of the contributions of this work.

The RESUN Framework

The two main components of RESUN are the evidential representation system and the control planner. The key feature of the evidential representation is its use of explicit, symbolic encodings of the *sources of uncertainty* (*SOUs*) in the evidence for the hypotheses. For example, a Track hypothesis in a vehicle monitoring system may be uncertain because its supporting sensor data might have alternative explanations as a Ghost or as part of a different Track or it may be uncertain because its evidence is incomplete or its correct Mission explanation is uncertain; these are possible sources of uncertainty for Track hypotheses. As interpretation inferences are made in RESUN, symbolic statements are attached to the hypotheses to represent their current sources of uncertainty. This allows the system to understand the *reasons* why its hypotheses are uncertain.

Control is provided by a script-based, incremental planner. A planning-based approach to control facilitates sophisticated control reasoning. The hierarchical goal/plan/subgoal structure created by a planner provides detailed context information as well as explicit decision options. This allows control reasoning to be very

[1] *Incremental hypothesize and test* (also known as *evidence aggregation*) means that to resolve uncertainty about a hypothesis, a system should attempt to locate all the data that would have been caused if the hypothesis were correct. *Differential diagnosis* means that a system should attempt to discount the possible *alternative explanations* for the hypothesis' supporting evidence.

[2] Both Hearsay-II and the planner of [Durfee & Lesser 1986] included techniques which implicitly did some limited differential diagnosis.

detailed; decision procedures can be highly context-specific and can reason explicitly about the choices. In addition, because planning-based control is inherently goal-directed, it can support active data gathering for applications like vehicle monitoring.

The main innovation in our planner is its *refocusing mechanism.* This mechanism can be used to postpone focusing decisions when there is insufficient information about the particular situation to make a conclusive choice. The ability to postpone focusing decisions results in a model of control in which there is not only a search for the correct interpretations, but also an explicit search for the best *methods* to use to pursue the interpretations. The refocusing mechanism also adds *opportunistic* control capabilities to the (goal-directed) planning mechanism by allowing focus points to change in response to a variety of events. This is crucial to the successful use of planning-based control. Interpretation requires data/event-directed control capabilities to deal with uncertainties in the data and problem solving knowledge as well as to handle dynamic situations (as in vehicle monitoring). The refocusing mechanism makes it clear that opportunism is not some special form of control that must be added to the planner, but that it simply results from particular types of conditions which should redirect the control search.

In the RESUN framework, interpretation is viewed as an incremental process of gathering evidence to resolve particular sources of uncertainty in the hypotheses. Control plans invoke actions that examine hypotheses and return information about the symbolic SOUs associated with the hypotheses. Focusing knowledge is applied to select the SOUs that will be used in further plan expansion to post goals to resolve uncertainty. These goals allow the system to identify methods that can resolve the current sources of uncertainty in the hypotheses. Focusing knowledge is again applied to select the best methods to pursue and the plans for these methods are refined to produce appropriate interpretation actions. This general process is repeated until the termination criteria are satisfied.

Termination is an important issues for interpretation. Interpretation systems not only must resolve uncertainty about the correctness of the hypotheses they create, they must also be sufficiently certain there are not additional answers which have not yet been identified. This is a critical issue because possible hypotheses are incrementally identified when doing constructive problem solving and it is typically infeasible to examine all of the data. As part of the RESUN approach, we have developed a high level model of the state of problem solving that is used to drive the overall interpretation process. This model represents the need to resolve uncertainty in existing hypotheses and to investigate the possibility of additional answers. For example, additional answers might be possible if some portion of the region of interest has not yet been examined by the system or if there is data which has not been ruled out as being able to support an answer. The problem solving model makes it possible to use goal-directed strategies to limit the amount of the data which is examined. This capability is important in applications like vehicle monitoring where there may be a number of sensors generating continuous streams of data.

Hypotheses and Sources of Uncertainty

The basis of the interpretation process is *abduction.* It is abductive inferences that identify possible explanations for data (and, conversely, possible *support* for hypotheses). Abductive inferences are uncertain due to the possibility of alternative explanations for the data. This is the basic underlying source of all interpretation uncertainty. However, our symbolic SOUs must represent more information than just the possible alternative explanations for hypotheses; there are several factors which influence the level of belief in hypotheses and thus several ways to go about resolving uncertainty.

Hypothesis correctness can only be guaranteed by doing complete differential diagnosis—i.e., discounting all of the possible explanations for the supporting data. Even if complete support can be found for a hypothesis there may still be alternative explanations for all of this support. However, while complete support cannot guarantee correctness, the amount of supporting evidence is often a significant factor when evaluating the belief in a hypothesis (this is the basis of hypothesize and test). For example, once a Track hypothesis is supported by correlated sensor data from a "reasonable" number of individual positions (Vehicle hypotheses), the belief in the Track will be fairly high regardless of whether alternative explanations for its supporting data are still possible. In addition, complete differential diagnosis is typically very difficult because it requires the enumeration of all of the possible interpretations which might include the supporting data—many of which may not be able to be conclusively discounted. Thus, a combination of hypothesize and test and (partial) discounting of critical alternative explanations must be used to gather sufficient evidence for interpretation hypotheses. Our representation of uncertainty is designed to drive this process.

Another important aspect of our evidential representation is its view of a hypothesis as a set of *extensions.* Each extension represents a different possible "version" of the hypothesis—i.e., a different binding for the hypothesis' parameters. The versions of a hypothesis that are of interest and must be represented are identified as part of the constructive problem solving process. A hypothesis' parameter values are constrained by the parameter values of its supporting data and hypotheses. Typically, evidence (especially incomplete evidence) will only partially constrain a hypothesis' parameters—i.e., there will be uncertainty about the correct values for the parameters. Thus, evidence for an interpretation hypothesis not only *justifies* the hypothesis, it can also *refine* it by further constraining its parameter values. However, because most interpretation evidence is uncertain, alternative sets of evidence may be pursued for a hypothesis. This produces multiple alternative versions of the hypothesis. In most blackboard systems, these versions are maintained as independent hypotheses; ignoring valuable information about the relationships between the versions. Using extensions, we can represent a high level of belief in a Track *hypothesis* (i.e., high belief that there is a vehicle moving through the monitored environment) despite great uncertainty about the correct version of the hypothesis (i.e., uncertainty about the correct path or identity of the vehicle).

Our model of interpretation uncertainty is based on the reasons why abductive inferences are uncertain, the factors that affect the belief in interpretation hypotheses, and our extensions model of hypotheses. The model specifies a set of SOU *classes* that characterize the uncertainty in interpretation hypotheses. These classes apply to all interpretation domains. In addition, we have identified a set of SOU class instances that are appropriate for particular applications. A discussion of SOU instances is beyond the scope of this paper (see [Carver 1990]).

Our model of interpretation uncertainty consists of the following SOU *classes* for hypothesis extensions (see [Carver 1990] for more detailed definitions):

partial evidence Denotes the fact that there is incomplete evidence for the hypothesis. For example, a Track hypothesis will have a *no explanation* SOU associated with it if it has not yet have been examined for valid Mission explanations and will have a *partial support* SOU if its supporting Vehicle hypotheses only cover a portion of the complete Track.

possible alternative support Denotes the possibility that there may be alternative evidence which could play the same role as a current piece of support evidence. This is an additional complication for differential diagnosis in interpretation problems as compared with classification problems.

possible alternative explanation Denotes the possibility that there may be particular alternative explanations for the hypothesis extension.

alternative extension Denotes the existence of a competing, alternative extension of the same hypothesis; using evidence which is inconsistent with other versions of the hypothesis. This is the primary representation of the relationships between hypotheses.

negative evidence Denotes the failure to be able to find some particular support or any valid explanations. Negative evidence is not conclusive because it also has sources of uncertainty associated with it—e.g., that sensors may miss some data.

uncertain constraint Denotes that a constraint associated with the inference could not be validated because of incomplete evidence or uncertain parameter values. This SOU represents uncertainty over the *validity* of an evidential inference; the other SOUs are concerned with the *correctness* of inferences.

uncertain evidence Technically, this is not another SOU *class*. *Uncertain support* and *uncertain explanation* SOUs serve as placeholders for the uncertainty in the evidence for a hypothesis because SOUs are not automatically propagated.

Figure 2 shows three extensions of a Track hypothesis along with their associated SOUs and parameters. Track-Ext_1 is an *intermediate extension* while Track-Ext_2 and Track-Ext_3 are alternative *maximal extensions*. The alternative extensions result from competing possible explanations of the Track as an Attack-Mission or as a Recon-Mission. This alternatives relationship between these Mission hypotheses is represented by the *alternative extension* SOUs in Track-Ext_2 and Track-Ext_3. These SOUs indicate that there is a negative evidential relationship between the extensions: more belief in Track-Ext_2 or Attack-Mission results in less belief in Track-Ext_3 or Recon-Mission (and vice versa). They also make it possible for the system to recognize that the uncertainty in Attack-Mission need not be directly resolved, but can be pursued by resolving the uncertainty in Recon-Mission or by resolving the uncertainty in the Track's parameter values (in order to limit its consistent interpretations). This example also demonstrates how extensions represent different versions of hypotheses: the uncertainty in the value of Track-Ext_1's ID parameter has been resolved differently by the alternative explanations. The uncertainty that results from each explanation only being consistent with a subset of the possible values for the Track's ID parameter is represented by *uncertain constraint* SOUs. These SOUs do not appear in the figure because they are maintained as part of the inferences; they are accessed through the "placeholder" *uncertain-explanation* SOUs which represent the overall uncertainty in the explanations.

RESUN's evidential representation system includes a scheme for numerically summarizing the symbolic SOUs using domain-specific evaluation functions. The summarization process produces a composite characterization of the uncertainty in a hypothesis in terms of an overall belief rating and the relative uncertainty contributions of the different classes of SOUs. This summarization is used in evaluating the satisfaction of termination criteria and when reasoning about control decisions. Having the composite rating allows for more detailed reasoning than would be possible with a single number rating. For example, it can distinguish between a hypothesis that has low belief due to a lack of evidence and one for which there is negative evidence. The composite rating also permits the use of modular evaluation functions (these evaluation functions effectively compute conditional probabilities—see [Pearl 1988]). Domain-specific evaluation functions are currently used because neither Bayes' Rule nor Dempster's Rule are generally applicable to interpretation due to

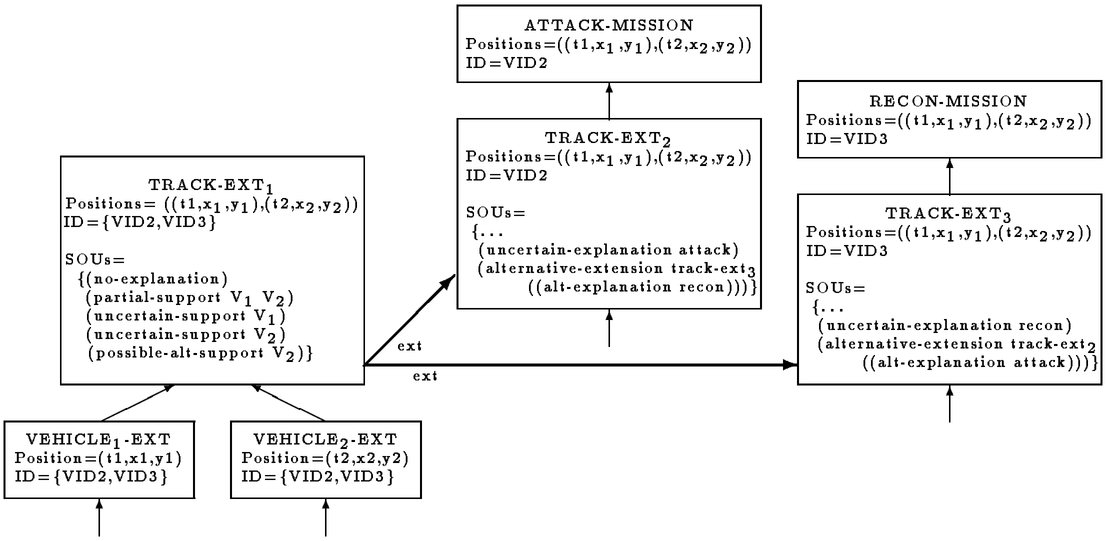

Figure 2: Example hypothesis extensions with their symbolic sources of uncertainty.

lack of independence of hypothesis evidence.

The RESUN representation of hypotheses and evidence addresses a problem that was first identified in Hearsay-II [Lesser & Erman 1977]: "The state information associated with a hypothesis is very local and does not adequately characterize the state(s) of the hypothesis network(s) connected to it ...the state information associated with an individual hypothesis must allow a KS to analyze quickly ...the role that the hypothesis plays in the larger context of the hypothesis networks it is part of." The representation of hypotheses as set of alternative extensions effectively maintains independent *contexts* that can be characterized by the summarization process.

Numeric representations of uncertainty like probabilities and Dempster-Shafer belief functions cannot to be used to identify methods for directly resolving uncertainties because they *summarize* the reasons why evidence is uncertain [Pearl 1988]. Our use of a symbolic representation of uncertainty is similar to [Cohen 1985]'s symbolic representations of the reasons to believe and disbelieve evidence which he calls *endorsements*. However, the work on endorsements did not produce any general formalism for representing and reasoning with symbolic evidence. Our representation is specific to abductive inferences and the needs of interpretation control.

Opportunistic Control Planning

The planner that we developed is a script-based, incremental planner. *Script-based planning* [Swartout 1988] means that the planning process is based on a set of control plan *schemas* that represent the possible methods that can be used to satisfy goals. Each non-primitive plan specifies a sequence of subgoals that implement the plan using a shuffle grammar that can express strict sequences, concurrency, alternatives, optional subsequences, and iterated subgoal subsequences. Each primitive plan represents an *action* that can be taken to immediately satisfy a goal. RESUN's format for specifying control plans is described in [Carver 1990].

Classical AI planners [Wilkins 1988] are not appropriate for domains like interpretation where the outcome of actions is uncertain and where external agents affect the world. We deal with these problems through *incremental planning* (interleaving planning and execution), allowing actions to return results, and by including explicit *information gathering actions*. Successful actions may return results that are bound to plan variables and influence further plan expansion. Information gathering actions allow the planner to maintain only that part of the world state which is needed and to make sure it is sufficiently up to date. *Data gathering actions* are similar to information gathering actions except that they are used to control active sensors.

As plans are refined and expanded, a structure like that shown in Figure 3 is created. Here the subgoal *Have-Ext-SOU*, the initial subgoal of the plan, *Eliminate-Extension-SOUs*, matches the primitive *Identify-Sources-of-Uncertainty*. When the primitive is executed, it returns a list of the SOUs in the specified hypothesis extension. This list is bound to the plan variable sou. Following this action, the plan is expanded further, posting the subgoal *Have-Eliminated-Ext-SOU*. This subgoal includes the *partial-support* binding of the variable sou which was selected through focusing. This subgoal matches two plans, one of which is selected by focusing for further refinement.

In a planning-based approach to control, control decisions—i.e., decisions about which domain actions to take next—result from a sequence of planner focusing decisions. Thus focusing heuristics represent *strategy* knowledge that selects the interpretation *methods*

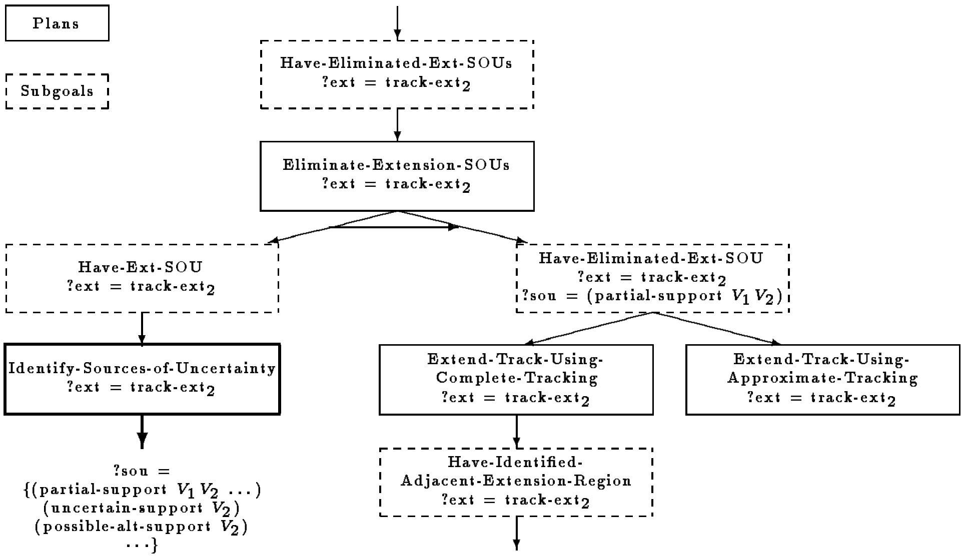

Figure 3: An example of the instantiated goal/plan/subgoal structure produced by the planner.

and method instances to be pursued. In RESUN, each focusing heuristic is associated with a particular control plan and can examine the instantiated planning structure. This provides detailed context information for decisions. Strategy knowledge is defined in terms of three classes of focusing heuristics. *Match* focusing heuristics select among competing control plans capable of satisfying a subgoal—i.e., competing methods. *Variable* focusing heuristics select among competing bindings for plan variables—i.e., competing method *instances. Subgoal* focusing heuristics select among the active subgoals for a plan instance when subgoals can be carried out concurrently, but it is preferable to sequence the subgoals (due to uncertainty over their ability to be satisfied, for instance).

The *refocusing mechanism* allows focusing heuristics to designate their decision points as *refocus points.* This is done by instantiating a *refocus unit* that specifies the decision point, the conditions under which refocusing should occur, and a refocus handler. When the refocus conditions are satisfied, the refocus handler is invoked and re-evaluates the choices made at the decision point—within the context of the further expanded plan. Using this mechanism, the system can deal with nondeterminism in focusing decisions by postponing decisions in order to gather more specific information about the situation. For example, when extending a Track hypothesis, the best direction to extend it in depends on the quality of the data which is actually available in each alternative region. The refocusing mechanism makes it possible to postpone the decision about where to extend the track until the plans for *both* alternative directions are expanded to a point where the relative quality of the data can be evaluated. When the plans have been expanded to this point, the decision is reconsidered and a the single best direction is selected to be pursued for the next track extension.

The refocusing mechanism also makes it possible to implement *opportunistic* control strategies that can shift the system's focus-of-attention between competing plans and goals in response to changes in the situation. This is possible because refocus units are evaluated and applied in a demon-like fashion and their conditions can refer to the characteristics of the developing plans and interpretations, and other factors such as data availability. For example, the amount of effort to be expended on one alternative can be limited or the arrival of critical data noted. Refocusing controls the system's backtracking since refocus points effectively define the backtrack points and the conditions under which the system backtracks. This provides the system with an intelligent form of nonchronological backtracking because it is directed by heuristic refocusing knowledge.

A number of planning-based control approaches have been developed, but none provide a completely suitable framework for interpretation driven by our symbolic SOUs. [Clancey 1986]'s tasks and meta-rules are really control plans and their substeps, but the framework is limited by the fact that meta-rules directly invoke subtasks so there is no ability to search for the best methods. The BB1 system [Hayes-Roth & Hewett 1988] has a different view of planning, in which plans select sequences of ratings functions rather than directly selecting actions. This limits its ability to support detailed, explicit control reasoning. Also, since BB1 relies on an agenda mechanism, it can be inefficient for interpretation problems involving large amounts of data because only a fraction of the possible actions will ever be taken (see [Hayes-Roth 1990] for recent work that

	experiment 1	experiment 2	experiment 3	experiment 4
plans instantiated	121	150	167	154
subgoals instantiated	309	428	462	440
hypotheses created	83	48	83	59
extensions created	216	116	213	152
inference actions	56	51	77	69
information actions	136	226	254	232
focusing decisions	82	93	108	97
planning time	11.6%	20.6%	14.4%	18.0%
focusing time	4.6%	7.0%	5.2%	6.3%
inference action time	81.1%	66.5%	76.2%	70.7%
information action time	2.6%	5.9%	4.3%	5.0%
total cpu time	31.3s	23.2s	36.3s	27.6s

experiment 1 Weak criteria for the acceptance of non-answers, no sophisticated strategies or differential diagnosis.
experiment 2 Like experiment 1, but sophisticated control strategies; still no differential diagnosis.
experiment 3 Like experiment 2, but with stronger criteria for acceptance of non-answers.
experiment 4 Like experiment 3, but using differential diagnosis strategies.

Figure 4: Some sample results from the experimental evaluation.

addresses this issue). The incremental planning approach of [Durfee & Lesser 1986] for a blackboard-based vehicle monitoring system is not a general planning mechanism. Its strategy of building abstract models of the data to guide the interpretation process is a particular problem-solving strategy that could be used in our system with the addition of appropriate abstraction actions. [Firby 1987]'s reactive planner uses a plan schema representation that is similar to ours, but does not address the issues of focusing the planner and obtaining and integrating knowledge about the current state of the world.

Status and Conclusions

In order to evaluate the RESUN framework, we have implemented the concepts with a simulated aircraft monitoring application. The implementation is in Common Lisp on a Texas Instruments Explorer using GBB [Gallagher, Corkill & Johnson 1988] to implement the hypothesis blackboard. Aircraft monitoring is a suitable domain for the evaluation because it has characteristics that exercise all of the capabilities of the system: there are large numbers of potential interpretations of the data due to the modeling of ghosting, noise, and sensor errors, there are complex interactions between competing hypotheses, and there can be multiple types of sensors some of which are active and controllable. The experimental results are presented and analyzed in [Carver 1990]. To date, the experiments have been desgined primarily to evaluate the usefulness of this framework for defining complex, context-specific interpretation strategies. We will discuss a few of the conclusions from the experimentation here.

One of the most important conclusions was confirmation that the combination of explicit control plans with context-specific focusing heuristics provides a flexible framework for developing sophisticated interpretation strategies. The modularity of the control plans and focusing heuristics as well as the ability to do explicit control reasoning makes it fairly easy to write and adapt control strategies. We also found that planning-based control is useful in a development environment because it makes it is apparent when additional strategies are required; with agenda-based blackboard control schemes, it can be difficult to determine whether adequate strategies have been defined (encoded in the scheduler rating function). While some flexibility may be lost with highly goal-directed control, we believe that a suitable set of strategies can be developed by testing a number of scenarios and by including some default method search strategies.

The results from a portion of one series of experiments are shown in Figure 4. These experiments are based on a data scenario in which there are two alternative tracks that was also used in [Durfee & Lesser 1986]. Experiment 1 used strategies that are comparable to those that would be found in conventional blackboard systems. For experiment 2, sophisticated goal-directed, context-specific control strategies were added—though the resolution of uncertainty was still based on hypothesize and test strategies. The results show that cpu-time reductions of 26%. were achieved and with more complex scenarios, reductions of up to 54% were acheived. These results are comparable to the performance improvements that were obtained in [Durfee & Lesser 1986] through the use of data abstraction and modeling techniques. Experiment 3 demonstrates how the system responds to changes in its goals. Here the level of certainty required to eliminate potential answers from consideration is increased over experiment 2. This forces the system to do additional work to disprove potential answers (the system is still not allowed to use differential diagnosis strategies). Experiment 4 demonstrates that the ability to use differential diagnosis strategies in resolving hypothesis uncertainty can result in substantial improvements in problem solving performance. Not only were cpu-time reductions of 24 to 28% achieved, but higher levels of confidence in solutions could also be obtained (there is a limit to the confidence that can be obtained with hypothesize and test alone).

Assessing control overhead from these experiments is problematic because the implementation has not yet been optimized and results depend on the relative costliness of the inference actions (which will vary from domain to domain.) Our inference functions are relatively simple; they do not contain substantial numeric calculations like Fast Fourier Transforms. Thus it is reasonable to expect lower overall overhead from other applications. Nonetheless, we compared figures for overhead with a study of BB1 [Garvey & Hayes-Roth 1989] and found overhead to be comparable.

In conclusion, this paper describes a new framework for building sensor interpretation systems. While most existing blackboard-based interpretation systems have been limited to using hypothesize and test strategies to resolve uncertainty, RESUN supports the use of more sophisticated strategies like differential diagnosis. The RESUN approach is based on a model that we developed of the uncertainty in abductive interpretation inferences. This model makes it possible to symbolically represent the sources of uncertainty in interpretation hypotheses. We also developed an incremental planner that can be used to implement methods which exploit this symbolic representation of uncertainty. The key innovation of the planner is its refocusing mechanism which makes it possible to handle nondeterminism in control decisions and adds opportunistic control capabilities to the goal-directed planner. We have found that the modularity of the control plans and the context-specific focusing heuristics provides a good framework for encoding complex control strategies. A detailed description and evaluation of the system is contained in [Carver 1990]. We are currently exploring the generality of RESUN using the domain of sound understanding in household environments [Lesser et al. 1991].

References

Carver, N., and Lesser, V. 1990. *Control for Interpretation: Planning to Resolve Uncertainty*, Technical Report 90-53, Computer and Information Science Department, University of Massachusetts.

Carver, N. 1990. *Sophisticated Control for Interpretation: Planning to Resolve Uncertainty.* Ph.D. diss., Computer and Information Science Department, University of Massachusetts.

Clancey, W. 1985. Heuristic Classification. *Artificial Intelligence* 27:289–350.

Clancey, W., and Bock, C. 1986. *Representing Control Knowledge as Abstract Tasks and Metarules*, Technical Report KSL 85-16, Knowledge Systems Laboratory, Computer Science Department, Stanford University.

Cohen, P. 1985. *Heuristic Reasoning About Uncertainty: An Artificial Intelligence Approach.* Pitman.

Durfee, E., and Lesser, V. 1986. Incremental Planning to Control a Blackboard-Based Problem Solver. *Proceedings of AAAI-86*, 58–64.

Erman, L., Hayes-Roth, F., Lesser, V., and Reddy, D. R. 1988. The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty. *Blackboard Systems.* Engelmore, R., and Morgan, T. eds. Addison-Wesley.

Firby, R. J. 1987. An Investigation into Reactive Planning in Complex Domains. *Proceedings of AAAI-87*, 202–206.

Gallagher, K., Corkill, D., and Johnson, P. 1988. *GBB Reference Manual*, Technical Report 88-66, Computer and Information Science Department, University of Massachusetts.

Garvey, A., and Hayes-Roth, B. 1989. An Empirical Analysis of Explicit vs. Implicit Control Architectures. *Blackboard Architectures and Applications.* Jagannathan, V., Dodhiawala, R., and Baum, L. eds. Academic Press.

Hayes-Roth, B. 1990. An Architecture for Adaptive Intelligent Systems. *Proceedings of DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control*, 422–432.

Hayes-Roth, B., and Hewett, M. 1988. BB1: An Implementation of the Blackboard Control Architecture. *Blackboard Systems.* Engelmore, R., and Morgan, T. eds. Addison-Wesley.

Lesser, V., and Erman, L. 1977. A Retrospective View the Hearsay-II Architecture. *Proceedings of IJCAI-77*, 790–800.

Lesser, V., and Corkill, D. 1983. The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem Solving Networks. *AI Magazine* (4)3:15–33.

Lesser, V., Nawab, H., et al. 1991. *Integrated Signal Processing and Signal Understanding*, Technical Report 91-34, Computer and Information Science Department, University of Massachusetts.

Nii, H. P., Feigenbaum, E., Anton, J., and Rockmore, A. 1982. Signal-to-Symbol Transformation: HASP/SIAP Case Study. *The AI Magazine* (3)1:23–35.

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

Peng, Y., and Reggia, J. 1986. Plausibility of Diagnostic Hypotheses: The Nature of Simplicity. *Proceedings of AAAI-86*, 140–145.

Swartout, W. ed. 1988. Report: DARPA Santa Cruz Workshop on Planning. *AI Magazine* (9)2:115–131.

Wilkins, D. 1988. *Practical Planning: Extending the Classical AI Planning Paradigm.* Morgan Kaufmann.

Williams, M. 1988. Hierarchical Multi-expert Signal Understanding. *Blackboard Systems.* Engelmore, R., and Morgan, T. eds. Addison-Wesley.