Control for Interpretation: Planning to Resolve Sources of Uncertainty*

Norman Carver and Victor Lesser

COINS Technical Report 90-53
July 1990

Abstract

Interpretation is a complex and uncertain process which requires sophisticated evidential reasoning and control schemes. We have developed an interpretation framework which models interpretation as a process of gathering evidence to resolve *particular sources of uncertainty* in the interpretation hypotheses. This allows us to directly resolve uncertainty through the use of *differential diagnosis* techniques instead of being limited to *evidence aggregation*. The key components of the approach are an evidential representation which includes explicit, symbolic statements of the sources of uncertainty in the evidence for the hypotheses and a script-based, incremental control planner. The control plan schemas which define the interpretation methods contain explicit *information gathering actions* which examine the symbolic sources of uncertainty associated with particular hypotheses. This information is used to post goals for resolving specific uncertainties in the current interpretations. These goals then direct the system to expand methods which are appropriate for resolving the uncertainties represented in the goals. Strategy knowledge is defined in focusing heuristics which are applied during the planning process to select the best methods and method instances to be pursued. Sometimes focusing decisions may not be able to be made definitively at the appropriate point during planning because there is insufficient information to select the most appropriate alternative. In these cases, the decisions may be postponed and multiple alternative methods partially expanded in order to accumulate sufficient information to select between them. This is accomplished with a *refocusing* mechanism which allows focusing decisions to be reconsidered once the alternatives have been expanded appropriately. Thus, the control process can be viewed as a search for the best methods to pursue as well as a search for the correct interpretations. The refocusing mechanism has also been extended to provide the goal-directed planner with a data-directed, opportunistic control capability. This framework is an alternative to a conventional blackboard system. It extends the blackboard representation of interpretation hypotheses and replaces the typical agenda-based blackboard control with a planner whose primitive plans correspond to knowledge source instantiations.

*This work was supported by the Office of Naval Research under University Research Initiative grant number N00014-86-K-0764 and DARPA contract N00014-89-J-1877.

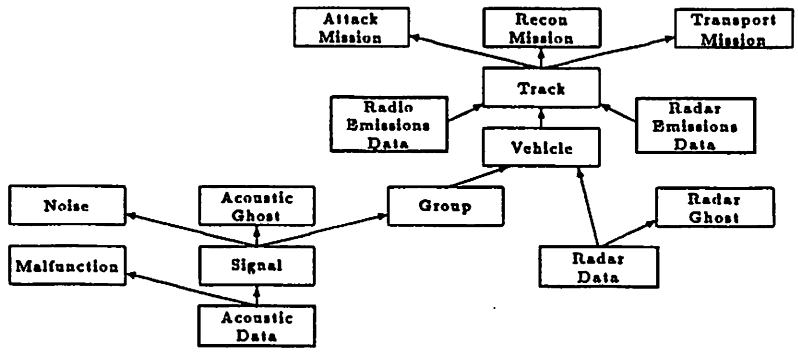

Figure 1: Interpretation type hierarchy for aircraft monitoring.

1 Introduction

Interpretation is the determination of a high-level, abstract view of sensor and other observational data. The interpretation process is based on a specification of the relations between the data and a set of abstraction types. The types form a hierarchy like the one in Figure 1 for aircraft monitoring. An interpretation system incrementally creates and extends *hypotheses* which represent instances of the various abstraction types by analyzing more data. Each of these hypotheses represents a possible *explanation* for some subset of the available data and, conversely, each hypothesis relies on the data or data abstractions for its *support*.

Interpretation can require sophisticated problem solving techniques because a combinatorial number of hypotheses may be able to be created to explain the data, creating each hypothesis may be computationally expensive, the correctness of the hypotheses will be uncertain, and a variety of methods may be used to resolve the uncertainty in each of them. The aircraft monitoring situation shown in Figure 2 illustrates the problem as it shows only a small portion of the potentially valid interpretations for even a simple set of sensor data. The large number of possible interpretations here is due to the multiple alternative interpretations for each piece of data, uncertainty in the data characteristics, weak constraints which limit how rapidly alternatives are excluded, and the possibility of missing data due to sensor malfunctions.

Existing interpretation systems (e.g., [1, 9, 16, 17, 24]) have limited strategies for resolving interpretation uncertainty because they don't understand the *reasons* why hypotheses are uncertain and because they don't understand their overall problem solving goals. This paper describes a new framework which we have implemented that makes it possible to define complex strategies for interpretation. In particular, it makes it possible to use *differential diagnosis* techniques to resolve interpretation uncertainty rather than being limited to *incremental hypothesis and test* as is the case for most interpretation systems. By incremental hypothesize and test (also called evidence aggregation) we mean that when a hypothesis has only partial support, the system should attempt to find the complete support which the hypothesis should have if it is correct. By differential diagnosis we mean that the system should attempt to discount the possible alternative explanations for a hypothesis' supporting data. Differential diagnosis has not been used explicitly in interpretation systems up to now. This is because of how difficult it can be to understand the relationships which develop between the hypotheses as alternative interpretations are created and extended. The framework we have developed provides explicit, symbolic representations of the *sources of uncertainty* in interpretation hypotheses and makes it possible to use this information to direct the interpretation process. A key contribution of the representation of uncertainty is that it allows the system to recognize the evidential relationships between alternative hypotheses.

Our framework is an alternative to conventional blackboard systems for interpretation. 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 [17]. However, explicit differential diagnosis techniques have not been exploited by blackboard-based interpretation systems due to the limitations of their evidence and control frameworks. HEARSAY-II did include a technique for implicitly doing some limited differential diagnosis through "KSI clustering." This involved pursuing sets of similarly rated hypotheses together. Though the similarly rated hypotheses were not necessarily competing alternative hypotheses, the technique was very useful when they were and caused no harm when they weren't. KSI clustering was developed because when similarly rated hypotheses were competing alternatives, HEARSAY's "island driving" strategy would cause whichever hypothesis was first extended to then be pursued to the exclusion of the other alternatives. This is because the system had no way to recognize that one alternative had become more highly rated than another simply because more evidence had been gathered for it—not because there was any evidence against its alternative. Even if such evidential relations had been represented they would have been difficult to exploit. In conventional agenda management, the rating of each KSI is done independently of rating the other KSIs. This local evaluation procedure limits the ability of the control system to understand the relationships among potential actions. Where explicit relationships among KSIs have been exploited [18], additional stages of processing beyond the local evaluation have been required.

In trying to add significant differential diagnosis capability, we have had to extend the representation of hypotheses and abandon the agenda control paradigm. Our representation of hypotheses maintains detailed information about the reasons hypotheses are uncertain and about the evidential relations between alternative hypotheses. We replaced the agenda-based control of standard blackboard systems with a planning mechanism where primitive plans correspond to an agenda-based blackboard system's knowledge source instantiations (KSIs). The planner provide detailed context for making control decisions and coordinating sequences of actions for meeting "long term" goals. For example, comparisons among alternative actions is facilitated by a planning approach to control because the subgoal structure of the instantiated plans identify the relationships among the actions. While our new control scheme does not maintain the modularity between knowledge sources and control as in a classic blackboard system, we feel that its advantages justify the change. We have also developed mechanisms which allow the goal-directed planner to retain the opportunistic control capabilities of more conventional blackboard systems.

Section 2 investigates the differences between interpretation and classification problems in order to make it clear why classification techniques alone do not suffice for interpretation. The interpretation framework that we have developed is introduced in Section 3. More detailed information about the evidential representation and the control planner are then given in Sections 4 and 5, respectively. The architecture of the resulting system is examined in Section 6. Section 7 describes the current status of our research and presents some results

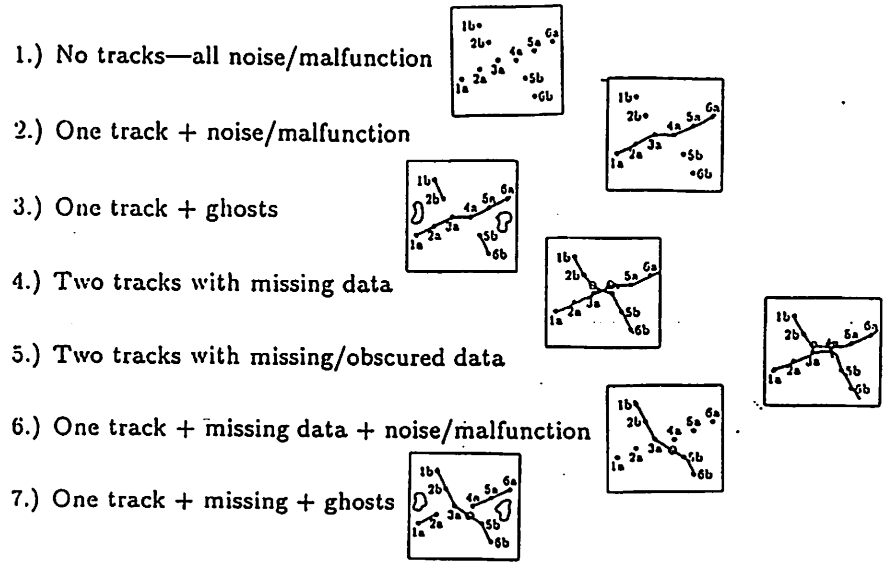

Figure 2: Combinatorial explosion of possible interpretations.

from the evaluation of the implemented system. Section 8 discusses related research on planning for control, Cohen's symbolic representation of evidence called endoresemnts, and Pearl's evidential reasoning techniques using Bayesian Networks. The paper concludes with a brief summary of the contributions of this research in Section 9.

2 Interpretation vs. Classification

In order to understand why interpretation is difficult, it is necessary to recognize the differences between interpretation and *classification* problems such as simple MYCIN-like medical diagnosis—e.g., [2, 22]. While interpretation involves classification, interpretation requires additional techniques beyond those which suffice for classification because of the combinatorics of the answer space. Clancey [5] distinguishes between *classification problem solving* where the problem solver *selects* the answer from among a set of pre-enumerated solutions and *constructive problem solving* in which the possible answers must be determined as part of problem solving. Clancey notes that constructive problem solving systems must be able to incrementally construct and extend hypotheses, maintain large numbers of incomplete hypotheses, and apply significant focusing mechanisms to control the construction process.

The enumeration of all potential hypotheses is precluded in interpretation problems because there are extremely large or even infinite numbers of hypotheses which must be considered. One reason for this is that interpretation hypotheses are compound structures which include parameters and some of these parameters typically have either large numbers of discrete values or even continuous values. Thus there are typically a very large or even infinite number of possible versions of each hypothesis. For example, Track hypotheses in an aircraft monitoring system include an "ID" parameter which identifies the type of the vehicle out of all possible aircraft types and also a "positions" parameter which gives the X–Y positions of the aircraft over time. The combinatorics of the representation of hypotheses are further complicated by the fact that each piece of evidence may support different subsets or ranges of the values for each parameter. This could require representations for all of the 2^n subsets of the hypothesis versions where n is the often very large number of possible hypothesis versions. For instance, each piece of acoustic sensor data is capable of ultimately supporting a number of different aircraft types because each piece of data represents a single acoustic signal frequency which could have been produced by a number of different aircraft. Uncertainty in the actual frequency being sensed leads to an even greater number of aircraft types which might be supported. It is only by combining many pieces of sensor data that this vehicle ID uncertainty can be resolved—i.e., that the subset of possible aircraft types which the available data supports can be determined.

An additional factor which precludes the pre-enumeration of all potential hypotheses is that there may be multiple correct instances of each type of hypothesis. For example, in an aircraft monitoring application, the number of aircraft which will be monitored in a given region over some time period is not known a priori. Thus the system must not only resolve its uncertainty about the correctness of any Track or Mission hypotheses that it creates, but must also be sure that it has examined enough of the data to create hypotheses for all possible aircraft. This means that the overall goal of an interpretation system is not only to resolve its uncertainty about the correctness of the hypotheses it has created, but to be sure that these hypotheses cover all of the valid interpretations. Another consequence of the possibility of multiple hypothesis instances is that it is never certain whether some hypothesis supports another hypothesis that it is capable of supporting or whether it actually supports a different *instance* of the other hypothesis—an instance which may not even be instantiated yet. For example, given some acoustic data which is able to support a particular Track hypothesis (through intervening Signal, Group, and Vehicle abstraction levels), it is possible that the data actually resulted from a different vehicle in the environment. In other words, the data should actually be used to support an alternative Track hypothesis—which may or may not have been created yet. Note that this source of uncertainty would remain even if all non-vehicle explanations for the data (e.g., ghost, noise, etc.) could be ruled out. This means that differential diagnosis in interpretation must consider the possibility of alternative explanation instances as well as the possibility of alternative explanation types as in classification.

When interpretation problems have been "solved" with classification techniques alone (see [19, 20]), what has actually been done is that many difficult aspects of the problems have been ignored. For example, continuous-valued hypothesis parameters have had their values partitioned into a very small set of discrete values. In the ship monitoring system in [20], position is represented as one of: {at-sea-zone1,...,in-channel,at-loading-dock,...}. While partitioning might be acceptable for some applications, it is not a general solution. Even if we acknowledge that sensors have limited resolution so the X-Y positions they return need not be represented by continuous values, many applications will still have too large a set of positions to be handled as discrete values. Another difficulty which has been ignored is the possibility of there being multiple instances of hypotheses. For example, [20] fuses several pieces of evidence about the position of *the ship* and even the position of *the ship* the next day without addressing the issue of how it is known that the pieces of evidence refer to the *same* ship. In data fusion terminology, [20] has ignored the issue of "correlation ambiguity," a major combinatorial problem [15].

3 Overview of the Framework

We have developed an interpretation framework which models interpretation as an incremental process of gathering evidence to resolve *particular sources of uncertainty* in the interpretation hypotheses. The two main components of the approach are the evidential representation system and the control planner. The key feature of the evidential representation is its use of explicit, symbolic statements of the *sources of uncertainty* (SOUs) in the evidence for the hypotheses. For example, a Track hypothesis in an aircraft monitoring system may be uncertain because its supporting sensor data might have alternative explanations as a ghost, a malfunction, or a different vehicle or it may be uncertain because its evidence is incomplete or its mission-level explanation is uncertain. As interpretation inferences are made in our system, symbolic statements are attached to the hypotheses to represent their current sources of uncertainty. Thus the Track hypothesis mentioned above might include statements that its supporting evidence is uncertain because of possible alternative explanations for the data, that its evidence is incomplete, and that there are alternative explanations for the Track. The use of the symbolic SOUs allows the system to understand the *reasons* why its hypotheses are uncertain so that uncertainty in the hypotheses can be resolved by more sophisticated methods than evidence aggregation. The evidential representation system also includes a method for summarizing the symbolic SOUs. This summarization produces a composite characterization of the uncertainty in a hypothesis' evidence in terms of the relative contributions of different classes of SOUs. The summarization is used by the system to evaluate the hypotheses with respect to the termination criteria and to identify the key uncertainties when making focusing decisions. As was discussed in Section 2, the termination criteria must include the goal of being sufficiently certain that hypotheses for all potential answers have been created as well as being sure of the created hypotheses. In order to evaluate this goal, we have included a high level model of problem solving which is used to drive the interpretation process. This model makes it possible for the system to recognize that further work is required to investigate the possibility of additional answers because no evidence has been generated for a region or because there are data points or hypotheses which have not yet been examined to see if they might support a new answer.

Control is provided by a script-based, incremental planner with heuristic focusing and refocusing mechanisms. A set of control plan schemas define the available interpretation *methods*. Each non-primitive control plan specifies a sequence of subgoals which accomplish the goal of the plan. Each *primitive* control plan represents an action which can immediately satisfy a goal and has a corresponding function (knowledge source) to execute the action. Besides the *inference actions* which create interpretation evidence, there are also *information gathering actions* and *data gathering actions.* Information gathering actions are explicitly included in the control plans where further plan expansion requires information about the state of the interpretations or the availability of data. The symbolic SOUs are integrated with the planning process through the use of information gathering actions. For example, a plan to resolve the uncertainty in a hypothesis would typically have an initial subgoal of identifying the current sources of uncertainty in the hypothesis. This subgoal is satisfied by a primitive plan which examines the hypothesis and returns a list of the symbolic SOUs associated with it. Further expansion of the plan to resolve uncertainty requires the application of heuristic focusing knowledge to select an SOU from this list to be resolved. A subgoal to resolve the

SOU is next posted and used to identify control plans (methods) which are able to resolve this particular SOU. Focusing heuristics then select the plan to be pursued and the new plan is instantiated and expanded. Data gathering actions are similar to information gathering actions, but are used to invoke active sensors to create new interpretation data. They can be used to tune the parameters of the sensors to produce the most appropriate kind of data for resolving the current uncertainties.

Planning for control facilitates explicit reasoning about control decisions because it provides detailed context for the decisions and because inference actions are selected through a series of choices among explicit alternatives. In our framework, control decisions are made by applying focusing heuristics during plan expansion to select the goals, plans, and plan variable bindings to be pursued. Thus, focusing heuristics represent *strategy* knowledge about the best methods and method instances for particular situations. Each focusing heuristic is associated with a particular control plan schema and is free to examine the control planning and interpretation structures when making a decision in order to determine the exact context for the decision. Sometimes definitive focusing decisions may not be able to be made at the point where goals, methods, or method instances must be selected because there is insufficient information about the particular situation. To deal with this, we have created a *refocusing* mechanism which can be used to postpone a focusing decision by allowing multiple alternatives to be partially expanded until sufficient information is developed to be able to make a decision. This technique might be used, for example, when doing tracking in an aircraft monitoring system which does not process all of the data in the sequence it is received—i.e., uses an opportunistic "island driving" approach to extending tracks. The best direction to extend an existing Track hypothesis 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 single best region is selected to be pursued for the Track extension.

The use of the refocusing mechanism to postpone focusing decisions results in a model of control in which there is not only a search for the correct interpretations, but also a search for the best methods and method instances to use to pursue the interpretations. While an agenda-based control scheme in a conventional blackboard system might be able to produce the same effect with a complex rating function, it would not be aware that it is performing this method search as our planner would be. In other words, while such a search may result from a rating function, the search process could not be represented explicitly—i.e., you could not directly express a control strategy like, "when deciding between method P and method Q, perform a search for the better method by partially expanding each one." In addition to its use in postponing focusing decisions, the refocusing mechanism is also used to add data-directed, opportunistic control to the goal-directed planning mechanism. Interpretation requires data-directed control as well as goal-directed control because in dynamic domains like aircraft monitoring, decisions which were correct when they were made may become incorrect as the situation evolves. To deal with this, the refocusing mechanism is used to reconsider decisions when specified conditions change—e.g., if new data becomes available in a critical region or if a plan is expanding in an unsatisfactory manner. This is done by posting a condition under which focusing decisions should be reconsidered. Refocusing allows the planner to be as opportunistic as an agenda-based control scheme, but again requires an explicit representation of the conditions for opportunism instead of leaving this to a rating function.

Most interpretation systems are limited to resolving uncertainty in their hypotheses through an *incremental hypothesize and test* strategy (also known as *evidence aggregation*). The framework we have described above makes it possible to define complex, explicit strategies for dealing with the combinatorics and uncertainty of interpretation. Since our representation identifies the reasons that hypotheses are uncertain, the system can pursue methods to directly resolve the causes of its uncertainty through *differential diagnosis* techniques. For example, the system can evaluate the possibility of alternative interpretations of the data such as those shown in Figure 2 using models of ghosting, noise, etc. Differential diagnosis does not entirely replace evidence aggregation, however, because it can be very expensive and because complete differential diagnosis may not be required to reach the desired level of confidence in hypotheses. Because the system does not use a fixed method for resolving its uncertainty the ability to define complex interpretation strategies is important: the system must be able to use the most appropriate method given its goals. The planning mechanism makes it possible for the actions taken by the system to be made very sensitive to the goals of the interpretation process (termination confidence criteria, time available, etc.) and the particulars of the situation (the characteristics of the data, the availability of sensors, *etc.*). Different goals and situations result in the system generating different subgoals, pursuing different methods for those subgoals, and doing different amounts of search for the optimal methods. This results in the creation of different subsets of the potential hypotheses and produces different levels of confidence in the results.

4 Hypotheses and Sources of Uncertainty

We have developed a model of interpretation hypotheses which symbolically represents the sources of uncertainty introduced by interpretation inferences, can produce a composite summary of the overall uncertainty in a hypothesis, and which maintains evidential relationships among hypotheses that are alternative explanations of the same data and among alternative versions of the hypotheses. This model is based on the requirements for control—i.e., that the system be able to understand the key sources of uncertainty in the current hypotheses so that it can identify the best methods for resolving its uncertainty.

An interpretation type hierarchy like that in Figure 1 is essentially a causal hierarchy which defines the instances of lower-level types which a type instance "causes." For example, an Attack Mission causes there to be a Track with certain parameter characteristics and a Track causes there to be a sequence of Vehicles with appropriate identity and position parameters which eventually causes there to be appropriate sensor data. Interpretation is then based on the notion of *abduction*: if B *causes* A then given an A, we might postulate a B as an *explanation* for A. Likewise, given a possible B, in order to provide *support* for B we will look for an A which B should have caused. Abductive inferences are uncertain inferences rather than logically correct inferences because of the possibility that there is some other cause of A—e.g., C. This is the basic underlying source of all interpretation uncertainty. However, our symbolic SOUs represent more information than just the possible alternative explanations for hypotheses because there are several factors which influence the belief in hypotheses and thus several ways to resolve uncertainty.

The two basic methods for resolving interpretation uncertainty that we want to direct the system to use are hypothesize and test and differential diagnosis (see Section 1). Hypothesis correctness can only be *guaranteed* by discounting all possible explanations for the supporting data—i.e., using differential diagnosis. However, this is typically difficult or even impossible because it requires the enumeration of all of the possible interpretations which might include the supporting data and because many of these alternatives may not be able to be conclusively discounted. Instead we use a combination of hypothesize and test and the discounting of critical alternative explanations to gather sufficient evidence for hypotheses. Hypothesize and test cannot guarantee the correctness of a hypothesis because even if complete support can be found for the hypothesis, there may still exist alternative explanations for all of this support. Despite this, the amount of supporting evidence which can be found for a hypothesis is often a significant factor when evaluating the belief in the hypothesis. For example, once a Track hypothesis has a significant number of Vehicle (position) hypotheses which support it, the belief in the Track will be quite high regardless of whether alternative explanations for its supporting data are still possible.

Another key aspect of the evidential representation is its view of a hypothesis as a set of *extensions* each representing a different version of the hypothesis. The need for the extension representation comes from the fact that it is infeasible to pre-enumerate and maintain all of the possible versions and subsets of versions of a hypothesis (see Section 2). Constructive problem solving means that we identify the versions of hypotheses which are of interest and which are to be represented as a part of the problem solving process. The values of a hypothesis' parameters are defined by the parameters of its support and explanation hypotheses. We handle parameter uncertainty by allowing the use of set and/or range values to represent the potentially correct values for a parameter and by allowing array values to be incompletely specified (again, see Section 2). For example, individual Vehicle hypotheses not only support Track hypotheses, they also define the ID of the vehicle and the positions of the Track over time. The value of the ID parameter of a Track hypothesis may be represented as a set of possible values due to uncertainty in the values of the ID parameters of its current supporting Vehicle hypotheses and the positions array will be incomplete based on the time span of the accumulated Vehicle support. Because of this parameter uncertainty, every time Vehicle evidence is added to a Track hypothesis it may affect the possible values of the ID and positions parameters of the Track hypothesis (see the discussion of Figure 3 below). In other words, gathering evidence for an interpretation hypothesis not only *justifies* the hypothesis, it also *refines* it by further constraining its parameter values. Because the connection between evidence and a hypothesis is uncertain due to the possibility of multiple instances of hypotheses (see Section 2), there may be alternative pieces of evidence which might be pursued—for the same hypothesis. Consequently, multiple versions of each hypothesis must be used to represent the alternative hypothesis refinements supported by different sets of (uncertain) evidence. When these versions are maintained as independent hypotheses, valuable information about the relationships between the versions is lost. Instead, we maintain alternative hypothesis versions as different extensions of a single root hypothesis. This allows us to understand, for example, that a Track *hypothesis* may very likely be correct even though we are still uncertain about the correct version or *extension* of the hypothesis—e.g., uncertain about the correct path over time. "A hypothesis" is defined by the initial evidence which creates the hypothesis and results in overall constraints on the parameter values of all extensions—e.g., a Track hypothesis would be constrained to some subset of the possible aircraft types for its ID and to some partial track of positions over time.

For each interpretation type T, the interpretation specification defines the type's support, $\mathbf{S}_T$, and its possible explanations, $\mathbf{E}_T$. The support, $\mathbf{S}_T$, is a set of support *sources*—i.e., $\mathbf{S}_T = \{\mathbf{S}_k\}$, Each support source consists of a set of type instance specifications—i.e., for each $\mathbf{S}_l \in \mathbf{S}_T$, $\mathbf{S}_l = \{S_{il}\}$ where each $S_{jl} \in \mathbf{S}_l$ is a type *instance* specification. By type instance specification, we mean an interpretation type along with parameter constraints—e.g., a Vehicle type with constraints on the position and ID parameters. For example, a Track hypothesis constrains the vehicle IDs of its supporting Vehicle hypotheses to be identical and their positions to be consistent with the movement characteristics of the particular vehicle. This definition of the support, $\mathbf{S}_T$, as a set of support sources, $\{\mathbf{S}_k\}$, is done to model domains like aircraft monitoring where there may be multiple *sources of evidence* for some types—e.g., a Vehicle hypothesis may be supported by radar data or by a set of Group hypotheses based on acoustic sensor data (see Figure 1). The possible explanations, $\mathbf{E}_T$, is a set of types, $\{E_i\}$, each of which might explain some type T hypothesis. For example, a Track hypothesis might be able to be explained as an Attack Mission, a Recon Mission, or a Transport Mission; these are the three possible explanation types for the Track type. Note, though, that because of the constraints which each of these Mission types places on the vehcle ID and positions of associated Tracks, each particular Track hypothesis may only be able to be explained by some subset of these Missions. Based on these definitions, every hypothesis of type T, H_T, is the result of a set of abductive inferences each of which is of the form: $\{H_{S_{jl}} \Rightarrow H_T\}$ (support evidence) or $H_T \Rightarrow H_{E_j}$ (explanation evidence) where $H_{S_{jl}}$ is a hypothesis corresponding to type instance specification $S_{jl} \in \mathbf{S}_l$, $\mathbf{S}_l \in \mathbf{S}_T$ and H_{E_j} is a hypothesis of type $E_j \in \mathbf{E}_T$.

Based on the factors which affect belief in hypotheses, the alternative methods for resolving uncertainty, and our extensions model of hypotheses, we have defined the following potential *classes* of SOUs for a hypothesis H_T:

- **partial evidence** – Denotes the fact that there is incomplete evidence for the hypothesis. For example, a No Explanation SOU means that no explanation has been determined and a Partial Support SOU means that for some support source, l, the current set of support hypotheses, $\{H_{S_{jl}}\}$ is incomplete—i.e., $\{S_{jl}\} \subset \mathbf{S}_l$ and $\{S_{jl}\} \neq \mathbf{S}_l$. For example, a Track hypothesis which has not yet have been examined for valid mission-level explanations will have an No Explanation SOU associated with it. typically have incomplete supporting Vehicle hypothesis evidence (no supporting Vehicle hypotheses for some times included within 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—i.e., that there exists a hypothesis $H'_{S_{jl}}$ which is the correct S_{jl} support rather than $H_{S_{jl}}$. This reflects the fact that though "a hypothesis" may be quite certain, there can still be uncertainty over the correctness of individual pieces of evidence for the hypothesis. This is an additional complication for differential diagnosis in interpretation problems as compared with classification problems (see Section 2). Classification problems do not have to contend with multiple instances of the types and so do not have this source of uncertainty. Interpretation must consider the possibility that there is alternative supporting evidence for a hypothesis—i.e., that there is a different *version* of the hypothesis which is actually correct.

- **possible alternative explanation** – Denotes the possibility that there may be alternative explanations for the hypothesis—i.e., that there exists a hypothesis H'_{E_k}, $E_k \in \mathbf{E}_T$ which is the correct explanation rather than H_{E_j}. These SOUs explicitly identify the possible explanation types based on the characteristics of the hypothesis.

- **alternative extension** – Denotes the existence of a competing, alternative extension of the same hypothesis. In other words, an alternative version of the hypothesis has been created using one or more pieces of evidence which are inconsistent with the existing versions of the hypothesis—i.e., using alternative support and/or an alternative explanation. This is one of the key representations of the relationships between hypotheses.

- **negative evidence** – Denotes the failure to be able to produce some particular support evidence, S_{jl}, or to find any valid explanations in $\mathbf{E}_T$. Negative evidence is not conclusive because it also has sources of uncertainty associated with it—e.g., that sensors may have missed 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 whereas the other SOUs are concerned with the *correctness* of inferences. See the example presented below for further explanation of the SOUs.

- **uncertain evidence** – Technically, this is not another source of uncertainty *class*. Uncertain evidence SOUs merely serve as placeholders for the uncertainty in the evidence for a hypothesis because the sources of uncertainty are not automatically propagated as evidential inferences are made. They denote the fact that an evidential inference is uncertain because the inference contains uncertain constraint SOUs and/or the hypothesis extension which is the basis for the inference contains SOUs.

Figure 3 shows three extensions of a Track hypothesis along with the associated symbolic SOUs. Track-Ext$_1$ is an intermediate extension which is being maintained while Track-Ext$_2$ and Track-Ext$_3$ are alternative *maximal* extensions of the Track hypothesis. The alternative extensions are the result of competing explanations of the Track as an Attack Mission or as a Recon Mission. This alternatives relationship is represented by the alt-extension SOUs in Track-Ext$_2$ and Track-Ext$_3$. These SOUs play an important role when evaluating the belief in these Track extensions (and the mission hypotheses they support) because they indicate uncertainty due to the existence of an alternative explanation—i.e., there is a negative evidential relationship between the extensions. The alt-extension SOUs can also be used to recognize that the uncertainty in the Attack Mission hypothesis, for example, need not be directly resolved, but can be pursued by resolving the uncertainty in the Recon Mission hypothesis or by resolving the uncertainty in the Track's parameter values in order to limit

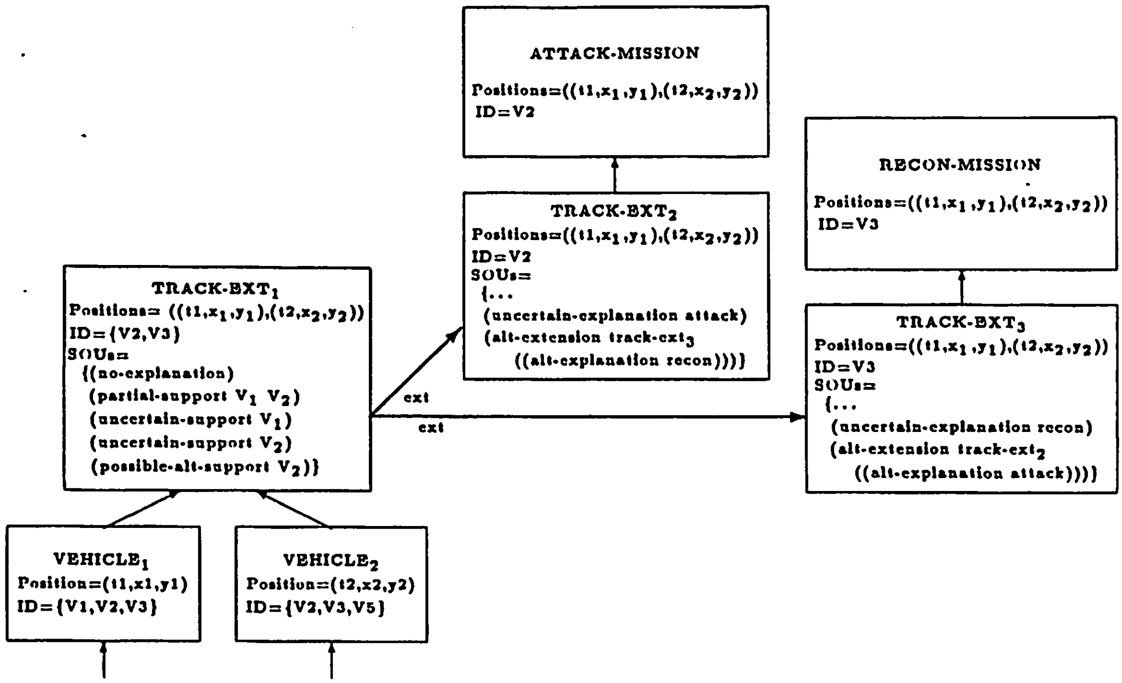

Figure 3: Example hypothesis extensions and sources of uncertainty.

its possible interpretations. Track-Ext$_2$ and Track-Ext$_3$ also contain placeholder uncertain-explanation SOUs which represent the fact that the explanations are uncertain. Here, this uncertainty includes the fact that each explanation is only consistent with a subset of the uncertainty ID values and so may not be valid. Uncertain contraint SOUs representing these uncertainties do not appear in the figure because they are actually maintained as part of the inferences and are not explicitly propagated to the extensions. This example also demonstrates how extensions represent different versions of hypotheses in terms of different parameter values: the uncertain values of ID in Track-Ext$_1$ have been resolved differently by the alternative explanations.

In addition to the symbolic uncertainty encoding, the system also provides a framework for producing a numeric characterization of the uncertainty in a hypothesis by summarizing the SOUs in the evidence hierarchy for the hypothesis. This composite numeric summary is used in evaluating termination criteria and in selecting the hypotheses to pursue and the methods to use to pursue them. The summarization process not only produces a probabilistic rating for a hypothesis, it also produces ratings for the different SOU classes contained in the evidence for the hypothesis: partial evidence, possible alternative explanations, possible alternative support, alternative extensions, negative evidence, and constraint uncertainty. This composite rating allows more detailed reasoning about termination and focusing decisions than would be possible with a single number rating. For example, it shows whether a hypothesis has low probability simply because little evidence has been gathered for it or whether there is evidence that it is incorrect. It can also show whether residual uncertainty results from actual competing hypotheses which may need to be examined or whether it is the result of the *possibility* of alternative evidence. The summarization process evaluates

PLAN DEFINITION:

Plan Name	Eliminate-Extension-SOUs
Description	Iteratively eliminates sources of uncertainty from the hypothesis extension ?ext while the belief in ?ext is less than ?belief.
Goal Form	(Have-Eliminated-Ext-SOUs ?ext ?belief)
Input Variables	(?ext ?belief)
Output Variables	()
Internal Variables	(?sou)
Grammar	(:ITERATION (< (BELIEF ?ext) ?belief) (:SEQUENCE Have-Ext-SOU Have-Eliminated-Ext-SOU))

SUBGOAL DEFINITION:

Subgoal Name	Have-Ext-SOU
Goal Form	(Have-Ext-SOU ?ext ?sou)
Input Variables	(ext)
Output Variables	(sou)

Figure 4: An example control plan definition.

the SOUs for a hypothesis using evaluation functions for each interpretation type which rate the relative contribution of each SOU to the uncertainty of the hypothesis and then use a combining function to produce the composite ratings for the hypothesis. The placeholder uncertain evidence SOUs are evaluated by evaluating the evidence they represent. This results in a recursive summarization process which examines the evidential structure supporting a hypothesis. Alternative Extension SOUs result in the evaluation of the alternative hypothesis extensions using the same process. The ratings of the relative contributions of the SOUs which are produced by the summarization process are maintained with the SOUs for use in focusing.

Our representation of hypotheses and evidence addresses a problem which was identified in HEARSAY-II [17]: "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 maintains independent versions of the evidence substructure which can be characterized by the summarization process.

5 Control Planning with Heuristic Focusing

The planning mechanism which we have developed to control our interpretation system was designed to maintain the benefits of planning systems (such as detailed context information and long-term coordination of actions) while being sufficiently reactive and opportunistic to cope with dynamic interpretation domains. One way we have made the planner more reactive is to make it script-based: a set of control plan schemas define the interpretation *methods* which are available to the system. The use of plan schemas is becoming popular in planning-based control schemes because it enhances the reactivity of a planner by limiting the "reasoning about time" which is one of the major sources of complexity in classical planners [14, 23]. Each non-primitive plan schema is defined using specifications like those in Figure 4. The goal which can be satisfied by the plan is specified by the Goal Form. The

```
Initialize current-focus-points to the top-level control plan instance
repeat: repeat: Pursue-Focus on each element of current-focus-points
                until null(current-focus-points)
        set current-focus-points to next-focus-points
        until null(next-focus-points)

Pursue-Focus(focus)
  case on type-of(focus):
    plan instance Focus on multiple-valued variable bindings to select plan instances.
                  Expand selected plan instances to next subgoals.
                  Focus on subgoals to select subgoals.
                  Match subgoals to control plans.
                  Focus on matching plans to select new plan instances for next-focus-points.
    primitive     Execute function associated with primitive to get status and results.
                  Update plans to select new focus element for next-focus-points:
                    propagate status and results of primitive to matching subgoal
                    and then up the control plan hierarchy to in-progress plan instance.
```

Figure 5: The basic control planning loop.

plan is realized by a sequence of subgoals which is defined by the Grammar clause. This clause uses a shuffle grammar to express strict sequences, concurrency, alternatives, optional subsequences, and iterated subsequences. The subgoal definitions for the plan identify the Goal Form, Input Variables, and Output Variables for each subgoal of the plan. These Goal Forms are used to identify control plans which can satisfy the subgoals by unifying the instantiated Goal Form with the Goal Forms of the set of defined control plans. Other clauses not shown in Figure 4 allow the specification of constraints on the variable values as well as precondition, satisfaction, and failure conditions to deal with subgoal interactions. These clauses are described in [4]. Primitive plans represent the *actions* that the system can take. Their specifications include the Goal Form, Input Variables, and Output Variables clauses and must also record the function (i.e., knowledge source) which implements the action.

The basic control planning loop is detailed in Figure 5 and an example control plan instance is represented in Figure 6. The plan instance shown is based on the plan defined in Figure 4. The initial expansion of the plan Eliminate-Extension-SOUs posts a single subgoal, Have-Ext-SOU. The Goal Form of this subgoal unifies with the Goal Form of the primitive plan Identify-Sources-of-Uncertainty. An instance of this primitive is created with its ?ext Input Variable bound by the unification process. This primitive returns a list of the SOUs in the specified hypothesis extension (actually it returns a multiple-valued value as explained below). The unification bindings result in the list of SOUs being bound to the ?sou variable of the subgoal and its containing plan. Further expansion of the plan results in posting the subgoal Have-Eliminated-Ext-SOU. However, since this subgoal has the variable ?sou as an Input Variable, focusing is applied to select the binding of ?sou to be used (again, see the discussion of multiple-valued values and variable focusing below). Two plans match the newly posted subgoal so focusing is applied to select the method to be pursued to satisfy the subgoal. The figure does not show examples of postponed focusing (see below). However, postponed focusing means that focusing on a binding for ?sou need not select a single binding, but may select multiple bindings to be (partially) pursued. In this case, an instance of the plan Eliminate-Extension-SOUs would be created for each of the alternative bindings of ?sou and each would be expanded. Likewise, focusing on a method

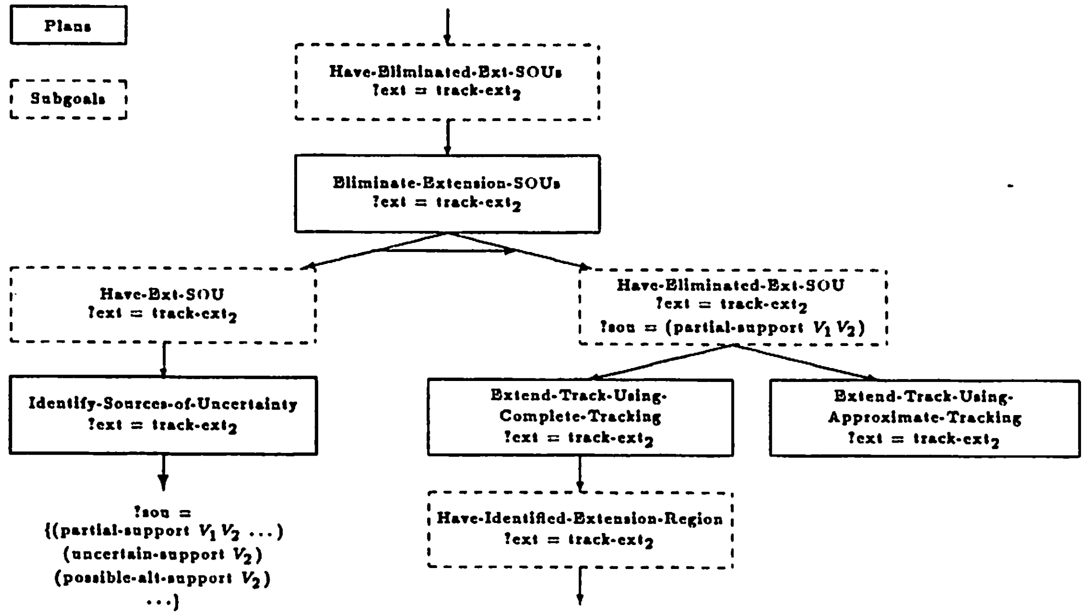

Figure 6: An example control plan instance.

to satisfy the subgoal Have-Eliminated-Ext-SOU need not select a single method.

Classical AI planning systems [23] form complete plans to solve their problems prior to executing them. However, forming complete plans prior to executing them is not appropriate for domains like interpretation where the outcome of actions is uncertain or where external agents can affect the world [14]. A related problem with classical planners is that they have to maintain complete world state models. This is very expensive and very difficult (it is the "frame problem") and is sometimes not possible for the same reasons that complete plans cannot be formed. We deal with these problems through incremental planning—interleaving planning and execution—and by including explicit *information gathering actions* in the plans. Incremental planning makes it possible to base the expansion of the plans on the outcome of earlier actions. One way that actions influence planning is through their execution status: actions may be considered to have succeeded or to have failed. A failed action will cause its matching subgoal to be considered failed unless the failure is handled via the refocusing mechanism (described below). Successful actions may also influence later plan expansion by returning results which are used to bind plan variables. Another response to the problems of classical planners is that we recognize three types of actions which may be included in the system: inference actions, information gathering actions, and data gathering actions. *Inference actions* correspond to knowledge source functions which create interpretation inferences and develop the interpretation hypotheses (see below for a further discussion of inference actions vs. knowledge sources). Inference actions are taken primarily for their side-effect of creating interpretation structure rather than for their results (although they may return results indicating the nature of the inference made). *Information gathering actions*, on the other hand, are taken solely for the results which they return. They are included in the control plan schemas at points where the system needs information about current state of the world—e.g., the state of the interpretations, the availability of data or sensors, etc. The symbolic sources of uncertainty are integrated with the control process in this way: an information gathering action is taken which returns a list of the current SOUs in a given hypothesis and this list is bound to a variable which will be used in a later subgoal. The use of information gathering actions allows the planner to maintain only that part of the world state which is necessary and to make sure it is sufficiently up to date. *Data gathering actions* instruct active sensors to gather additional data. These actions can help to resolve uncertainty because they can be used to tune the parameters and the positioning of appropriate sensors. The ability to actively direct the gathering of data requires that a system understand what it is that it wants to do. This understanding is provided in our system via the symbolic SOUs and the subgoals in the planning structure.

While we have stated that our inference actions correspond the knowledge sources (KSs) in a blackboard system, the inference functions are not identical to the functions that would be used to implement standard blackboard system KSs. Inference functions do not include the precondition component of blackboard KSs and they must determine the appropriate SOUs to be associated with the hypotheses they modify. Each KS in a conventional blackboard system has two major components: a precondition and an action. Our inference functions only include what would be the action portion of a KS. The precondition portion of a KS determines the applicability of the KS and creates bindings of KS variables to be used in instantiating KSs (creating knowledge source instantiations (KSIs)). In our system the functionality of the precondition portion of a KS is handled by the planning process: inference actions will only be selected when they can satisfy subgoals that planning has determined are appropriate. Binding of variables in inference function occurs via unification of the primitive plans representing the inferences with the matching subgoals. The other difference between inference actions and KSs is that in addition to creating and modifying hypotheses, inference actions must also determine the SOUs which will be associated with the hypotheses. In other words, KS actions must be extended to determine the additional information about the sources of uncertainty in the inferences they create.

The planning process is controlled by a set of *focusing heuristics*. In general, there will be many partial control plan instances which could be further elaborated at any point in the processing. These plan instances represent alternatives of pursuing different hypotheses, resolving different SOUs, and using different methods. One of the major control issues for interpretation is the development of an effective focus-of-attention scheme. In our system focusing is accomplished as part of the process of refining and expanding the control plans through the selection of methods and method instances to be expanded. Heuristic focusing knowledge represents meta-level *strategy knowledge* relative to the method knowledge represented by the plan schemas. In other words, while the plan schemas define the valid methods for accomplishing goals, the focusing heuristics represent strategy knowledge for selecting the *best* methods and method instances to pursue in particular situations. One of the key reasons for using a planning approach to control is to provide detailed context information for the decisions. Each focusing heuristic is associated with a particular control plan schema and is free to examine the hierarchy of control plan instances to determine the goals and subgoals which are affected by a decision. This eliminates the problem of conflicting heuristics which can occur with global meta-level focusing heuristics like Davis' meta-rules [8]. Instead of having general heuristics like "if the data has x characteristics then prefer it .6," we define heuristics which identify the best type of data for each specific context—i.e., for each particular purpose where the purpose is identified by the planning structure. For example, the choice of using radar data or acoustic sensor data would depend on the source of uncertainty the data is to be used to resolve as would the relative importance of data characteristics such as time slice data density or loudness (of acoustic data).

Control decisions must be made at a number of points in the planning process so we partition the focusing knowledge into three different classes of heuristics: match, variable, and subgoal. *Match* focusing heuristics select among competing methods when there are multiple control plans capable of satisfying a subgoal (multiple control plan schemas whose Goal Forms unify with a subgoal Goal Form). Each match heuristic is defined in association with a specific subgoal of a particular control plan schema. *Variable* focusing heuristics select among competing bindings for plan variables—i.e., they select among competing plan instances. Competing variable bindings are represented as *multiple-valued values.* Actions return multiple-valued values when focusing is to be invoked to select a plan parameter value. This is, in effect, a representation of uncertainty about the correct binding for a parameter. The use of multiple-valued values is shown in Figure 6. Here, the primitive Identify-Sources-of-Uncertainty doesn't return just a list of the current SOUs, but rather returns a multiple-valued value for the variable ?sou. The matching subgoal, Have-Ext-SOU, stands for the goal of having the SOU to be eliminated in the plan. Thus the use of a multiple-valued value represents the uncertainty over the best SOU to be eliminated. This then causes focusing to be applied to select the binding of ?sou to be pursued in the next subgoal, Have-Eliminated-Ext-SOU. The multiple-valued value framework is used so that all focusing is done as a meta-level process relative to the basic planning process. The alternative would be to include specific focusing subgoals within the plans, but this would place these focusing actions in a different context from the other focusing actions and would make it impossible to use the refocusing mechanism discussed below. Each variable heuristic is defined in association with specific a set of variables (since multiple variables may have to be focused at a single decision point) of a particular control plan schema. *Subgoal* focusing heuristics select among the active subgoals for a plan instance when the control plan grammar specifies that certain sets of subgoals may be satisfied in parallel. Despite the ability to be pursued in parallel, it may be preferable to sequence the subgoals because of uncertainty over their ability to be satisfied. Rather than encode such heuristic knowledge in the plans by sequencing the subgoals in the plan grammar, subgoal focusing may be used to sequence the subgoals without forcing this to always be done and by taking the particulars of the situation into account in the sequencing.

The refocusing mechanism allows focusing heuristics to designate their decision points as *refocus points.* This is done by instantiating a *refocus form* which specifies the decision point, the condition for refocusing, and a refocus handler. When the refocus condition is satisfied, the refocus handler is invoked and reevaluates the choices made at the decision point—within the new context. This allows the system to deal with the nondeterminism in focusing decisions by postponing the decisions or by making preliminary decisions. Postponing a decision means that multiple paths are partially expanded in order to gather more specific information about the situation before selecting the best alternative. This requires a refocus form which specifies the conditions under which the alternatives should be reevaluated and how to reevaluate them. Preliminary focus decisions are used when one alternative is likely

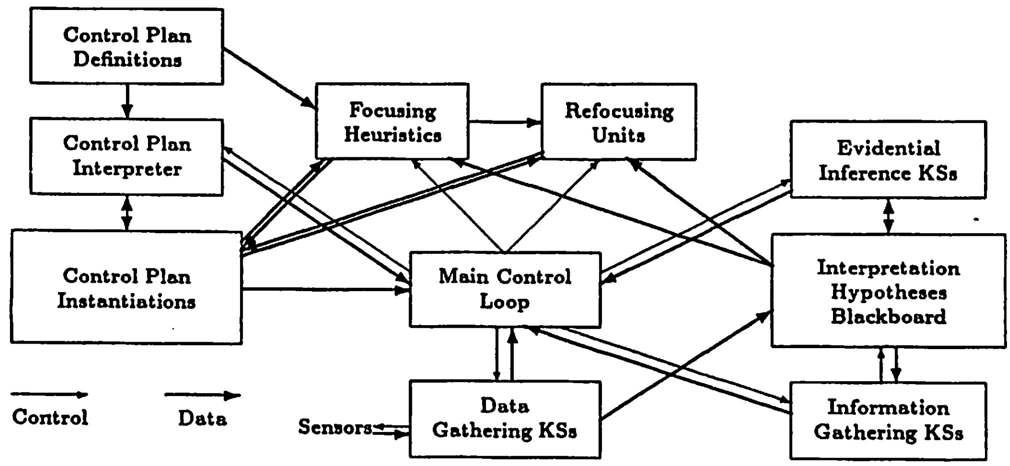

Figure 7: The architecture of the system.

to be the best—subject to certain reservations about its progress or under a particular assumption about the situation. The refocus conditions can then monitor the progress of the choice or the validity of the assumptions. For example, this can be used to limit the amount of effort expended on one alternative by setting a limit on the amount of time to be expended. Refocus forms 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. This makes it possible to define opportunistic strategies which shift the system's focus-of-attention between competing plans and goals in response to changes in the situation. 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.

6 Architecture

Figure 7 illustrates the relations between the major components of our system. The Main Control Loop module is responsible for executing the basic control loop of the system (see Figure 5). The Control Plan Interpreter uses the Control Plan Definitions to create, expand, and update the Control Plan Instantiations. Interpretation hypotheses are created by executing primitive plans which are implemented by corresponding Evidential Inference KSs. The interpretation hypotheses are maintained on the Interpretation Hypotheses Blackboard along with the symbolic SOUs. Non-inference primitive plans (actions) correspond to particular Information Gathering KSs or Data Gathering KSs. The Focusing Heuristics are associated with particular Control Plan Definitions. They are applied as appropriate by the Main Control Loop module to focus the planning process. When they are executed, Focusing Heuristics may also create Refocusing Units which can modify the flow of control under particular conditions by refocusing within the Control Plan Instantiations. The condition components of the Refocusing Units may refer to the state of the Control Plan Instantiations or the Interpretation Hypotheses Blackboard. Refocusing Units are effectively handled as demons.

7 Status

In order to evaluate this framework, we have implemented the concepts and simulated an aircraft monitoring application. The implementation is in Common Lisp on a Texas Instruments Explorer using GBB [11] to implement the data and hypothesis blackboards. Interpretation methods are defined in terms of 40 control plans consisting of approximately 150 subgoals and more than 40 primitives (actions). Aircraft monitoring is a suitable domain for the evaluation because it has complex characteristics which exercise all of the capabilities of the system. For example, multiple sources of evidence are available through the simulation of multiple types of sensors, some of these sensors are active and under the control of the system, there are complex interactions between competing hypotheses, and there are large numbers of potential interpretations of the data due to the modeling of ghosting, noise, and sensor errors. We believe that a system which is able to successfully handle a variety of problems within the aircraft monitoring domain will be able to handle a range of interpretation problems. The experiments which we have run have been designed to evaluate the ability to define complex interpretation strategies using the framework. Success has been measured by demonstrating:

- The system's actions are *responsive* to changes in the termination criteria, the characteristics of the data, and the a priori likelihoods for the sources of uncertainty.
- Performance is improved using *differential diagnosis* methods as compared to restricting the system to evidence aggregation.
- Performance can be improved by defining *complex heuristic strategies* using context-specific focusing heuristics, an explicit search for best methods via postponed focusing decisions, and opportunistic control using refocusing.

One series of experiments was run with a number of variations in the control configuration using the data scenario in Figure 2. The results of some of these experiments are shown in Figure 8. These four experiments are in "batch" mode with all of the data being available to the system prior to the start of the interpretation process. In the four scenarios, the focusing heuristics select the data at times 3 or 4 for the creation of an initial possible answer (Track) hypothesis because of the lower density of data at these times. Because of this, the system initially constructs the "a" Track shown in alternative 2 of Figure 2. As part of its goal of resolving uncertainty over whether there are additional possible answers, the system then attempts to construct a Track hypothesis using the uninterpreted "b" data. This second Track cannot be completed without assuming that there is missing data because it otherwise requires data that is supporting the first Track (which is strongly believed given the particular data in these scenarios). In Scenario 1, with criteria which requires relatively weak evidence to rule out possible answers, the fact that the second Track cannot be completed without assuming missing data is sufficient (i.e., the system does not gather evidence to discriminate between alternatives 2 and 3 of Figure 2). Scenario 2 demonstrates how the refocusing mechanism can be used to improve system performance. As compared with Scenario 1, the system now postpones its decisions about how to extend Tracks until it partially examines the available data. Experiments with the refocusing mechanism demonstrate that it can

	scenario 1	scenario 2	scenario 3	scenario 4
plans created	134	147	194	151
subgoals created	428	443	548	455
inference actions	68	53	79	65
information actions	227	244	275	250
focusing decisions	74	80	104	84
refocusing decisions	0	5	7	5
planning time	4.6s	5.0s	6.2s	5.3s
focusing time	1.2s	1.3s	1.7s	1.5s
refocusing time	0s	.17s	.23s	.17s
inference time	9.6s	7.5s	10.8s	8.9s
information time	1.1s	1.2s	1.6s	1.3s
total time	16.5s	15.2s	20.5	17.2s

scenario 1 Weak acceptance criteria for non-answers, no refocusing, no differential diagnosis.
scenario 2 Like scenario 1, but with refocusing.
scenario 3 Like scenario 2, but with strong acceptance criteria for non-answers.
scenario 4 Like scenario 3, but using differential diagnosis to resolve non-answer uncertainty.

Figure 8: Some example results from the experimental evaluation of the system.

reduce the number of interpretation inferences which the system performs, but there is often a tradeoff in terms of increased control reasoning. Despite this, because of the relative costs of control and inferencing, refocusing results in an 8% decrease in the execution time for Scenario 2 as compared with Scenario 1. Scenarios 3 and 4 demonstrate the responsiveness of the system to changes in the termination criteria and also demonstrate the advantages of using differential diagnosis. By increasing the strength of the acceptance criteria for ruling out possible answers, the system is forced to gather more evidence to rule out the "b" Track. In Scenario 3 this can only be done by more completely examining the data because the control plans which encode the differential diagnosis methods have been disabled. Fewer inferences and plans are generated in Scenario 4 because the system is able to use methods to directly investigate whether the "b" data is due to ghosting or noise (the data in this scenario is consistent with a ghost of the "a" Track). Here, differential diagnosis results in a 16% reduction in total execution time.

It is highly problematic to assess the absolute "control overhead" of our framework from these experiments. One reason is that the experiments are based on an implementation which has been designed more for research flexibility than for efficiency. While the inference functions (KSs) are relatively straightforward and could be little improved, the planner could be made significantly faster. This brings up the issue of the complexity of the inference functions. The more complex the inference functions, the lower the control overhead. Our inference functions do not contain substantial numeric calculations (e.g., Fast Fourier Transform) nor do the hypotheses contain large numbers of parameters or parameters with complex structures. Thus it is reasonable to expect lower overall overhead from more realistic inference functions. For example, in experiments conducted with PROTEAN [12], the control overhead ranged from only 8 to 24% of the total problem solving time due to the complexity of the inference KSs being used. If the relative cost of the inference functions were to double then control overhead in scenario 1 drops from 42% of the total execution time to 26% and it drops to 13% if the relative cost of the inferences increases by a factor of five. Likewise, the reduction in execution time between scenarios 1 and 2 increases from 8% to 13% when relative inference costs are doubled and to 18% when they increase by a factor of 5. Finally, it must be noted when comparing these results with those of an agenda-based blackboard system that the our inference functions do not contain the precondition component of a standard blackboard KS. Instead, this functionality is subsumed by the planner (see Section 5). In other words, not all of what is listed as planning time and information time in our performance data is actually "overhead" in the sense that it is additional work relative to the interpretation process as it would be realized by an agenda-based blackboard system. The complete set of experiments are described in [4].

8 Related Research

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 that the evidence is uncertain [21]. We have chosen to maintain a representation of the reasons that hypotheses are uncertain in order to allow the system to use a wide range of strategies to resolve uncertainty. Our use of a symbolic representation of uncertainty is similar to Cohen's [7] symbolic representations of the reasons to believe and disbelieve evidence which he calls *endorsements*. However, whereas Cohen was trying to develop a semantics for general evidential reasoning, our representation is tailored to the needs of interpretation control. This means that we only had to be concerned with the one type of well-understood inference which is the basis for interpretation: abductive inference. Cohen's task was very difficult because he was trying to capture the uncertainty in a wide variety of poorly understood types of inferences and develop methods for combining the endorsements from these inferences. Because of the complexity of the problem, the work on endorsements did not result in any general-purpose formalism for representing and reasoning with symbolic uncertainties. Our work demonstrates that it is possible to maintain and reason with detailed information about the sources of uncertainty in evidence when dealing with well-defined types of inferences.

Of course, a numeric summarization of the uncertainty in the evidence is still required in order to evaluate the hypotheses with respect to the termination criteria and to make focusing decisions. We have developed a summarization process which produces probability ratings and composite characterizations of the classes of SOUs for hypotheses and their extensions. Since an interpretation specification is effectively a specification of causal relations, the network of interpretation hypotheses which are constructed by the interpretation process are similar to a Bayesian network [21]. Since a Bayesian network along with Pearl's propagation scheme is an important new method for computing with probabilistic information, it is useful to understand how our evidential representation relates. As we have discussed already, it is impossible to pre-enumerate the possible versions of interpretation hypotheses or even the set of possible hypotheses; hypotheses and the versions of interest must be constructed as appropriate to the evidence which is examined. This means that interpretation is not simply a matter of instantiating a belief network [22] and then propagating probability information as new evidence is added to update the probability distributions for the network variables (where hypotheses would be represented as "multi-valued variables"). The key issue is, again, the need to *construct* new hypotheses and hypothesis versions as evidence is gathered. In other words, adding evidence will require changes to the network and even the creation of new networks—not just the propagation of probabilities. Despite these differences, we meet one of the main conceptual goals that Pearl has stated [21]: "to make intensional systems [1] operational by making relevance relationships explicit." Like Pearl, we use a propagation approach which does not suffer from the problems of extensional approaches. Our evaluation functions compute the appropriate *conditional probabilities* for hypotheses because they use the SOUs to identify the evidence that is relevant to the hypothesis and can recognize when relevant information has been changed. Thus our system includes one of the key features of the Bayesian network formalism: the use of the causal links to recognize information which is relevant to a hypothesis and thereby limit what must be considered out all the known evidence. We currently initiate the summarization and propagation process only on demand—that is, when the level of belief in a hypothesis is needed. Speed is not as critical an issue here since the size of interpretation structures is small relative to the kind of knowledge bases that Pearl is concerned with. We continue to do research to see whether Pearl's autonomous propagation approach could be adapted to our representation to produce more efficient revisions. It is important to recognize, however, that while Pearl's work addresses the issue of evaluating belief given a set of evidence, it does not address the problem of identifying evidence which could be gathered to resolve uncertainty—unless nodes for such evidence are already instantiated in the network (without any "diagnostic evidence"). Thus Pearl's work does not eliminate the need for explicit representations of the factors affecting belief if one is to decide how to resolve uncertainty.

A planning approach to control is advantageous for use with our symbolic representation of uncertainty since a planner develops a model of the system goals and the methods being used to pursue these goals. There are a number of systems which use a planning approach to control, but none of these provide a completely suitable framework for interpretation. Clancey's tasks and meta-rules [6] are really control plans and their substeps, but the framework is limited by the fact that meta-rules directly invoke subtasks. This means that there is no ability to encode strategies that can perform a search to determine which alternative methods to pursue. Hayes-Roth's work on blackboard control systems [13] has a different view of planning in which plans select sequences of ratings functions rather than directly selecting actions (as in our planner). The control blackboard approach relies on an agenda mechanism which considers (rates) all the currently possible actions on each loop. An agenda provides a high degree of opportunism, but can be inefficient for problems like interpretation where only a small fraction of the possible actions will be taken. Focusing as part of the planning process as we do, provides, in effect, a partitioned agenda: only those

[1]Intensional systems are those which correctly represent the context-sensitive nature of conditional probabilities. For example, the conditional probability $P(C \mid A)$ is *not* equivalent to a rule that says "anytime A is known then conclude C with the specified level of belief." Conditioning on A actually means "anytime A is known *and nothing else relevant is known ....*" Thus if, say, B is also known, then the appropriate conditional probability to use might be $P(C \mid A, B)$ rather than $P(C \mid A)$. Systems which do not properly reflect context-sensitivity (e.g., rule-based systems) are referred to as extensional or parallel certainty inference approaches.

actions immediately applicable to the in-focus plan or goal are identified and considered. Opportunism is maintained by the refocusing mechanism where the phenomena which should trigger the reconsideration of decisions are explicitly stated. Our control model makes it possible to reason explicitly about focusing and method search decisions because it makes a series of local decisions instead of making complex global decisions which are mediated by implicit numeric ratings as is done in the control blackboard approach because of its reliance on an agenda mechanism. The incremental planning approach of Durfee and Lesser [9] for a blackboard-based vehicle monitoring system builds abstract models of the data and uses these models to guide processing. The use of an abstract model of the data is one type of problem-solving strategy which could be accomodated in our system with the addition of appropriate abstraction operators. Both the data model and the planner are closely tied to vehicle tracking, though, so their application to general interpretation is limited. Firby's work on a reactive planner [10] uses a framework for representing and instantiating plans which is similar to ours. However, this work does not address the critical issue of focusing the planning process and does not provide a method for determining the state of the world and integrating this knowledge with the planning process.

9 Summary

This paper describes a new framework for building interpretation systems for complex tasks. These tasks are potentially combinatorially explosive in terms of the number of interpretations which can be generated and the number of alternative approaches which can be exploited. Conventional interpretation systems have limited models of the state of the interpretation process which limits their control strategies to some variant of incremental hypothesize and test. More sophisticated strategies such as differential diagnosis have not been exploited in interpretation because of the difficulty of understanding the relationships among hypotheses which result as alternative interpretations are incrementally constructed.

Our framework is based on a view of interpretation as an incremental process of gathering evidence to resolve particular sources of uncertainty in the hypotheses. We symbolically represent the sources of uncertainty resulting from interpretation inferences. This allows us to understand the reasons why hypotheses are uncertain, to characterize the reasons in a meaningful way, and to maintain evidential relationships among hypotheses which are alternative explanations for the same data. We have developed an incremental control planner which exploits this symbolic representation of uncertainty to identify methods which can directly resolve the current uncertainties in the interpretations. This makes it possible to use a differential diagnosis approach instead of being limited to evidence aggregation.

The framework represents an alternative to conventional blackboard systems. The representation of hypotheses has been extended to include explicit information about the reasons that hypotheses are uncertain. Agenda-based control has been replaced with a planner which provides context for making control decisions and can coordinate sequences of actions. Necessary opportunism is maintained in the goal-directed planner via the refocusing mechanism. A detailed description and evaluation of the system is contained in [4].

References

[1] Azarewicz, Jerome, Glenn Fala, Ralph Fink, and Christof Heithecker, "Plan Recognition for Airborne Tactical Decision Making," *Proceedings of AAAI-86*, 805–811, 1986.

[2] Buchanan, Bruce and Edward Shortliffe, editors, *Rule-Based Expert Systems*, Addison-Wesley, 1984.

[3] Carver, Norman, Victor Lesser, and Daniel McCue, "Focusing in Plan Recognition," *Proceedings of AAAI-84*, 1984, 42–48.

[4] Carver, Norman, *Sophisticated Control for Interpretation: Planning to Resolve Uncertainty*, Ph.D. Thesis, Computer and Information Science Department, University of Massachusetts, 1990.

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

[6] Clancey, William and Conrad Bock, *Representing Control Knowledge as Abstract Tasks and Metarules*, Technical Report KSL 85-16, Knowledge Systems Laboratory, Computer Science Department, Stanford University, 1986.

[7] Cohen, Paul, *Heuristic Reasoning About Uncertainty: An Artificial Intelligence Approach*, Pitman, 1985.

[8] Davis, Randall, "Meta-Rules: Reasoning about Control," *Artificial Intelligence,* volume 15, 179–222, 1980.

[9] Durfee, Edmund and Victor Lesser, "Incremental Planning to Control a Blackboard-Based Problem Solver," *Proceedings of AAAI-86*, 58–64, 1986.

[10] Firby, R. James, "An Investigation into Reactive Planning in Complex Domains," *Proceedings of AAAI-87*, 202–206, 1987.

[11] Gallagher, Kevin, Daniel Corkill, and Philip Johnson, *GBB Reference Manual*, Technical Report 88-66, Computer and Information Science Department, University of Massachusetts, 1988.

[12] Garvey, Alan and Barbara Hayes-Roth, "An Empirical Analysis of Explicit vs. Implicit Control Architectures," in *Blackboard Architectures and Applications*, V. Jagannathan, Rajendra Dodhiawala, and Lawrence Baum, editors, Academic Press, 1989.

[13] Hayes-Roth, Barbara, "A Blackboard Architecture for Control," *Artificial Intelligence*, 26, 1985, 251–321.

[14] Hendler, James, Austin Tate, and Mark Drummond, "AI Planning: Systems and Techniques," *AI Magazine*, vol. 11, no. 2, 61–77, 1990.

[15] Lakin, W.L., J.A.H. Miles, and C.D. Byrne, "Intelligent Data Fusion for Naval Command and Control," in *Blackboard Systems*, Robert Engelmore and Tony Morgan, editors, Addison-Wesley, 1988.

[16] Laskowski, Sharon and Emily Hofman, "Script-Based Reasoning for Situation Monitoring," *Proceedings of AAAI-87*, 819–823, 1987.

[17] Lesser, Victor and Lee Erman, "A Retrospective View the HEARSAY-II Architecture," *Proceedings of IJCAI-77*, 790–800, 1977.

[18] Lesser, Victor, Daniel Corkill, Robert Whitehair, and Joseph Hernandez, "Focus of Control Through Goal Relationships," *Proceedings of IJCAI-89*, 497–503, 1989.

[19] Lowrance, John and Thomas Garvey, *Evidential Reasoning: An Implementation for Multisensor Integration*, Technical Note 307, SRI International, 1983.

[20] Lowrance, John, Thomas Garvey, and Thomas Strat, "A Framework for Evidential-Reasoning Systems," *Proceedings of AAAI-86*, 896–903, 1986.

[21] Pearl, Judea, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufman, 1988.

[22] Peng, Yun and James Reggia, "Plausibility of Diagnositic Hypotheses: The Nature of Simplicity," *Proceedings of AAAI-86*, 140–145, 1986.

[23] Wilkins, David, *Practial Planning: Extending the Classical AI Planning Paradigm*, Morgan Kaufman, 1988.

[24] Williams, Mark, "Hierarchical Multi-expert Signal Understanding," in *Blackboard Systems*, Robert Engelmore and Tony Morgan, editors, Addison-Wesley, 1988.