24

Another Look at Frames

David E. Smith and Jan E. Clayton

The success of MYCIN-like systems has demonstrated that for many diagnostic tasks expert behavior can be successfully captured in simple goal-directed production systems. However, even for this class of problems, difficulties have arisen with both the representation and control mechanisms. One such system, PUFF (Kunz et al., 1978), has established a creditable record in the domain of pulmonary function diagnosis. The representation problems in PUFF are manifest in a number of rules that have awkward premises and conclusions. The control problems are somewhat more severe. Physicians have criticized PUFF on the grounds that it asks questions that do not follow a logical line of reasoning and that it does not notice data that are atypical or erroneous for the determined diagnosis.

In the CENTAUR system, described in Chapter 23, an attempt was made to correct representational deficiencies by using prototypes (frames) to characterize some of the system's knowledge. A more complex control scheme was also introduced. It made use of *triggering rules* for suggesting and ordering system goals, and included an additional attention-focusing mechanism by using frames as an index into the set of relevant rules.

In an attempt to carry the work of Aikins one step further, we have constructed an experimental system for pulmonary function diagnosis, called WHEEZE. Our objectives were to provide a uniform declarative representation for the domain knowledge and to permit additional control flexibility beyond that offered by PUFF or CENTAUR. To achieve the first of these objectives, all of PUFF's rules have been translated into a frame representation (discussed in Section 24.1). The second objective, control flexibility, is achieved by using an agenda-based control scheme (discussed

---

This chapter is an expanded version of a paper originally appearing in *Proceedings of the First National Conference on Artificial Intelligence*, Stanford, Calif., August 1980, pp. 154–156. Used with permission of the American Association for Artificial Intelligence.

in Section 24.2). New goals for the agenda are suggested by the success or failure of other goals on the agenda. In the final section, results and the possibilities of generalization are discussed.

24.1 Representation

24.1.1 The Language

We have chosen to use a representation language called RLL (Greiner and Lenat, 1980). The language is *frame-based,* where a frame consists of a set of *slots,* or attributes. We did not rely on the special features of RLL in any fundamental way. Any of the multitude of frame-based languages would have served equally well.

24.1.2 Vocabulary

In our knowledge base, there are three different kinds of frames that contain domain-specific diagnostic knowledge and knowledge about the case: assertion frames, patient frames, and patient datum frames.

Assertion Frames

The majority of the diagnostic knowledge is captured in a set of frames called *assertions.* Most assertions in the knowledge base are about the physiological state of the patient, e.g., "the patient's total lung capacity is high." But there are other types of assertions as well, such as "the total lung capacity measurement is erroneous." The organization of an assertion frame is shown in Figure 24-1.

An assertion may be related to other assertions in the knowledge base in several ways as shown in Figure 24-1. The substantiating evidence for an assertion is specified in the Manifestation slot for the assertion. This slot can be thought of as a set of links to secondary assertions that contribute to the confirmation of the assertion in question. It has been necessary to allow a considerable richness of combinations of manifestations for an assertion; consequently, each entry in the slot may be an individual manifestation or a simple function of individual manifestations, such as OneOf, TwoOf, TwoOrMoreOf, SomeOf, etc. Associated with each manifestation link is a number indicating the *importance* of the link in suggesting belief or disbelief in the assertion. The ManifestationOf slot is the inverse

Isa	Assertion
Description	<commentary>
Manifestation	<a list of assertions on which this assertion depends>
ManifestationOf	<a list of assertions that this assertion is a manifestation of—the inverse of the Manifestation slot>
Certainty	<a number between −1000 and 1000 that indicates to what degree the assertion is believed, if its manifestations are believed>
SuggestiveOf	<related assertions that are worth investigating if this assertion is believed>
ComplementaryTo	<related assertions that are worth investigating if this assertion is not believed>
CategorizationOf	<the patient datum that this assertion is concerned with>
CategoryCriterion	<the allowed range of the patient datum corresponding to this assertion>
DegreeOfBelief	<a number between −1000 and 1000 that indicates to what degree the assertion is believed>
Findings	<text to be reported to the user if this assertion is believed>

FIGURE 24-1 Organization of an assertion frame.

of the Manifestation slot; i.e., it contains a list of the assertions that have that assertion as a manifestation.

The Certainty slot, in WHEEZE, is an indicator of how likely an assertion is, given that its manifestations are believed. If the manifestations are strong indicators of the assertion, the Certainty slot will have a high value. The Certainty slot is a property of the knowledge rather than a statement about a particular consultation.

When an assertion is directly related to a patient datum, it is termed a *categorization* of that patient datum. This relationship is specified by the CategorizationOf and CategoryCriterion slots of the assertion. CategorizationOf indicates which patient datum the assertion depends on, while CategoryCriterion specifies the range in which the value must be for the assertion to be verified. For example, the assertion "the patient's TLC is greater than 110" (TLC stands for total lung capacity) would be a categorization of the TLC value with the category criterion being value>110.

The relationship may also be used in the reverse manner. A high-level datum such as SeverityOfDisease could be defined as one of a disjoint set of assertions being true (MildDisease, ModerateDisease, etc.), in which case the categorization relationship might be used to determine the datum from the assertions.

Each assertion has a DegreeOfBelief slot associated with it indicating to what degree the assertion is believed to be true in that particular consultation. The value of this slot can be any integer between −1000 and 1000, where 1000 indicates complete faith and −1000 means total denial of the assertion. It may also take on the value Unknown, indicating that the knowledge needed to determine the degree of belief of the assertion is not known. Note that there is a distinction made between a degree of belief that has not yet been investigated, a degree of belief that has been investigated but cannot be determined due to insufficient evidence (degree of belief Unknown) and a degree of belief that indicates equal positive and negative evidence (DegreeOfBelief = 0).

Unlike the Certainty slot, the DegreeOfBelief is determined by the system during the consultation. For an assertion that has only the categorization relationship (no manifestations), the DegreeOfBelief depends only on the Certainty of the assertion and on the patient datum being in the specified range. For assertions *with* manifestations, the DegreeOfBelief of the assertion can be a general function of the Certainty of the assertion, the DegreeOfBelief of each of its manifestations, and the importance attributed to each manifestation. The function used in MYCIN and PUFF is a simple thresholding mechanism, where, if the minimum of the antecedents is above some threshold (generally 200), the DegreeOfBelief is effectively set to the certainty factor. Importance measures provide additional flexibility by permitting the antecedents of a rule to be weighted. Several different combination mechanisms have been considered:

1. Sum the products of the DegreeOfBelief slots and the importance factors for each manifestation, then use a thresholding mechanism.
2. Sum the products of the DegreeOfBelief slots and the importance factors for each manifestation, then multiply this by the certainty factor.
3. Threshold the minimum of the DegreeOfBelief/importance ratios for the manifestations.

There are two assertion slots that indicate related assertions worth pursuing when an assertion is confirmed or denied. The SuggestiveOf slot contains a list of assertions to investigate if the current assertion is confirmed. Conversely, the ComplementaryTo slot is a list of assertions that should be pursued if the current assertion is denied. These slots function like the "triggering" rules in CENTAUR since they suggest goals to investigate.

The Findings slot of an assertion contains text that should be printed out if the assertion is confirmed. In PUFF, this text was contained in the conclusion portions of rules.

Isa	Patient
Age	<the patient's age>
Sex	<the patient's sex>
PackYearsSmoked	<the number of cigarette-smoking years specified in number of packs per day times number of years of smoking>
TLC	<the value of the total lung capacity for the patient>
RDX	<the referral diagnosis>
ConfirmedAssertions	<assertions that have already been confirmed for this patient>
DeniedAssertions	<assertions that have a DegreeOfBelief less than 0>
Agenda	<a pointer to an agenda frame containing assertions worth pursuing>

FIGURE 24-2 Organization of a patient frame.

Patient Frames

Information about the patient is kept in a frame named after that patient. In general, it contains slots for all of the patient data and for the state of the consultation. As shown in Figure 24-2, the majority of the slots in the patient frame contain the values of test data, derived data, or more general facts about the patient. Most of these values are entered directly by the physician; however, there are data that are derived or calculated from other values. The slots in the patient frame do not contain any information about obtaining the value for that slot. Instead, that information is kept in the corresponding patient datum frame (discussed below). The Confirmed-Assertions and DeniedAssertions slots keep track of the assertions that have already been tested. The Agenda slot contains a pointer to the agenda frame for the patient. It is important to note that the patient frame does not contain any heuristic knowledge about the system. Its only purpose is to hold current information about the patient.

Patient Datum Frames

In addition to patient and assertion frames, there are frames in the knowledge base for each type of patient datum (as shown in Figure 24-3). These frames indicate how a datum is obtained (whether it is requested from the physician or derived from other data), what a typical value for the datum

Isa	PatientDatum
Description	<commentary on this specific datum>
ToGetValue	<how to get the value of this datum if it is not known>
Categorization	<the set of assertions that are categorizations of this datum>
TypicalValue	<the value of this datum expected for a normal patient>

FIGURE 24-3 Organization of a patient datum frame.

might be, and what categories the value may be placed in. When the value of a patient datum is requested and not yet known, the frame for that patient datum is consulted and the information about how to obtain that datum is applied. This information takes the form of a procedure in the ToGetValue slot of the frame.

For a given patient datum, there may be many low-level assertions that are categorizations of the datum. These are specified by the Categorization slot. For example, the Categorization slot of TLC (total lung capacity) might contain the assertions TLC=80to100, TLC=100to120, TLC<80, and TLC>120, indicating that there are four major categories of the values. Thus the patient datum contains heuristic knowledge about how the datum is derived and how it relates to assertions in the network.

24.1.3 Translation

The process of translating a PUFF rule into a WHEEZE assertion consists of several steps. First, an assertion must be created embodying the conclusion and findings of the rule. Next, assertions corresponding to each of the antecedents of the rule must be constructed (if they are not already present) and added to the Manifestation slot of the assertion. If a manifestation is a categorization of some patient datum, then the CategorizationOf and CategoryCriterion slots for that manifestation must be filled in accordingly, and the frame describing that patient datum must be created.

Figure 24-4 is an example of how a particular PUFF rule was translated into our representation. The conclusion of the rule corresponds to the assertion and findings. The antecedents became the manifestations of the assertion. Quite often the manifesting assertions are not already present in the knowledge base and must be created. For example, the assertion frame RDX-Asthma (meaning "referral diagnosis of asthma") had to be added to the knowledge base when the RefractoryAsthma frame was created, since it is one of the manifestations of RefractoryAsthma. The patient

PUFF Rule 42

If: 1) There are postbronchodilation test results, and

2) The degree of reversibility of airway obstruction of the patient is less than or equal to slight, and

3) Asthma is one of the referral diagnoses of the patient

Then: It is definite (1000) that the following is one of the conclusion statements about this interpretation: The poor response to bronchodilators is an indication of an asthmatic condition in a refractory state.

REFRACTORY-ASTHMA

Isa	PhysiologicalState
Manifestation	(OAD BronchodilationTestResults RDX-Asthma (*OneOf OADReversibility-None OADReversibility-Slight))
Certainty	1000
DegreeOfBelief	
Findings	The poor response to bronchodilators is an indication of an asthmatic condition in a refractory state.
ComplementaryTo	((RefractoryAsthma-None 5))

FIGURE 24-4 PUFF rule and corresponding WHEEZE frame for refractory asthma.

datum RDX (referral diagnosis) also had to be added, since RDX-Asthma was specified as a categorization of RDX. Most of the other rules in the system were translated in an analogous fashion.

While there is not a one-to-one mapping between the representations we have used and the rules in PUFF, we can imagine automating the process. The most difficult problem in conversion is to create meaningful and consistent names for the assertions in the knowledge base. In most cases we used some combination of keywords in the conclusion of the rule we were mapping into the assertion (as in Figure 24-4).

24.2 Control Structure

Depth-first, goal-directed search is often used in production systems because questions asked by the system are focused on specific topics. Thus the system appears to follow a coherent line of reasoning, more closely mimicking that of human diagnosticians. There are, however, many widely recognized limitations. No mechanism is provided for dynamically select-

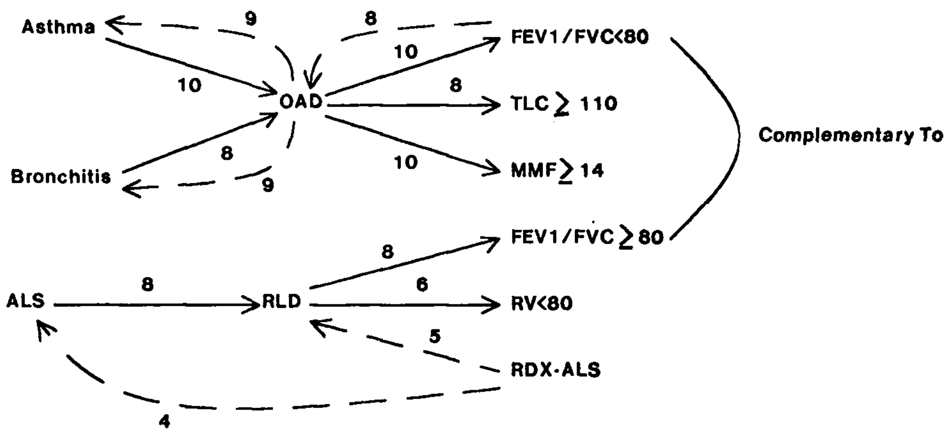

FIGURE 24-5 A simplified portion of the WHEEZE knowledge base. The solid lines indicate Manifestation links (e.g., OAD is a manifestation of Asthma); the dashed lines represent SuggestiveOf links. The numbers represent the corresponding importance and SuggestiveOf values of the links. (Key: ALS = amyotrophic lateral sclerosis; FEV1 = forced expiratory volume at one minute; FVC = forced vital capacity; MMF = maximal midexpiratory flow; OAD = obstructive airways disease; RDX = referral diagnosis; RLD = restrictive lung disease; RV = residual volume; TLC = total lung capacity.)

ing or ordering the initial set of goals. Consequently, the system may explore many "red herrings" and ask irrelevant questions before encountering a good hypothesis. In addition, a startling piece of evidence (strongly suggesting a different hypothesis) cannot cause suspension of the current investigation and pursuit of the alternative.

For the assertion network in Figure 24-5, a depth-first, goal-directed system like PUFF would start with the goals Asthma, Bronchitis, and ALS (amyotrophic lateral sclerosis) and work backwards in a goal-directed fashion toward OAD (obstructive airways disease) and RLD (restrictive lung disease) and then toward FEV1/FVC<80, MMF≥14, etc. In contrast, the CENTAUR system would make use of triggering rules to allow primitive data (e.g., RDX-ALS and FEV1/FVC<80) to suggest whether ALS and OAD were worth investigating and the order in which to investigate them. It would then proceed in a goal-directed fashion to try to verify those goals.

Expert diagnosticians use more than simple goal-directed reasoning. They seem to work by alternately constructing and verifying hypotheses, corresponding to a mix of data- and goal-directed search. They expect expert systems to reason in an analogous manner. It is therefore necessary that the system designer have some control over the reasoning behavior of the system. These intuitions, and the work on triggering described in Chapter 23, have led us to adopt a control mechanism that permits a combination of backward chaining and forward (data-driven) exploration together with any search strategy ranging from pure depth-first to pure breadth-first search. This control structure is implemented by using an agenda, with each suggested assertion being placed on the agenda according to some specified priority. The control strategy is as follows:

1. Examine the top assertion on the agenda.
2. If its subassertions (manifestations) are known, the relative belief of the assertion is determined. If confirmed, any assertions of which it is suggestive are placed on the agenda according to the specified measure of suggestivity. If denied, complementary assertions are placed on the agenda according to their measures of suggestivity.
3. If it cannot be immediately verified or rejected, then its unknown manifestations are placed on the agenda according to their measures of importance and the agenda level of the original assertion.

By varying the importance factors, SuggestiveOf values, and the initial items placed on the agenda, numerous control strategies are possible. For example, if high-level goals are placed on the agenda initially and subgoals are always placed at the top of the agenda, depth-first, goal-directed behavior will result. Alternatively, if low-level data are placed on the agenda initially and assertions suggested by these data assertions are always placed below them on the agenda, breadth-first, data-driven behavior will result. More commonly, what is desired is a mixture of the two, in which assertions suggest others as being likely and goal-directed verification is employed to investigate the likely assertions. The example below illustrates how this can be done.

In the knowledge base of Figure 24-5, suppose that RDX-ALS is confirmed, suggesting RLD to the agenda at level 5 and ALS at level 4. RLD is then examined, and since its manifestations are unknown, they are placed at the specified level on the agenda. The agenda now contains FEV1/FVC≥80 at level 8, RV<80 and RLD at level 5, and ALS at level 4. FEV1/FVC≥80 is therefore selected. Suppose that it is found to be false. Its complementary assertion (FEV1/FVC<80) is placed at level 8 on the agenda and is immediately investigated. It is, of course, true, causing OAD to be placed at level 8 on the agenda. The diagnosis proceeds by investigating the manifestations of OAD; and, if OAD is confirmed, Asthma and Bronchitis are investigated.

Although many subtleties have been glossed over in this example, it is important to note that:

1. The manipulation of SuggestiveOf and importance values can change the order in which assertions are examined, therefore changing the

order in which questions are asked and results are printed out. (In the example, FEV1/FVC was asked for before RV.)

2. Surprise values (data contrary to the hypothesis currently being investigated) may suggest goals to the agenda that are high enough to cause suspension of the current investigation. (The surprise FEV1/FVC value caused suspension of the RLD investigation in favor of the OAD investigation. If the suggestivity of the link from FEV1/FVC<80 to OAD were not as high, this would not have occurred.)
3. Low-level data assertions cause the suggestion of high-level goals, thus selecting and ordering goals to avoid irrelevant questions. (In the example, RLD and ALS were suggested and ordered by the low-level assertion RDX-ALS.)

24.3 Conclusions

It is no surprise that WHEEZE exhibits the same diagnostic behavior as its predecessors, PUFF and CENTAUR, on a standard set of ten patient test cases. The three systems are also roughly comparable in efficiency. WHEEZE and CENTAUR are somewhat slower than PUFF, but this may be misleading, since little effort has been expended on optimizing either of these systems.

The frame representation described in Section 24.1 has proved entirely adequate for capturing the domain knowledge of both PUFF and CENTAUR. In some cases, several rules were collapsed into a single assertion frame. In other cases, intermediate assertions, corresponding to common groups of clauses in rule premises, were added to the knowledge base. This had the effect of simplifying other assertion frames. The combination of representation and control structure also eliminated the need for many awkward interdependent rules and eliminated the need for screening clauses in others.

There are several less tangible effects of using a frame representation. Our purely subjective view is that a uniform, declarative representation is often more perspicuous. As an example, all of the interconnections between assertions about disease states are made explicit by the Manifestation and ManifestationOf slots. As a result, it is easier to find all other assertions related to a given assertion. This in turn makes it somewhat easier to understand and predict the control flow of the system.

Since the agenda-based control mechanism includes backward-chaining and goal-triggering capabilities, it has also proved adequate for capturing the control flow of PUFF and CENTAUR. In addition, the flexibility of agenda-based control was used to advantage. Suggestiveness and importance factors were used to change the order in which questions were asked and conclusions printed out. They were also used to eliminate the need to order carefully sets of antecedent assertions.

There is evidence that mixed goal-directed and data-directed control models human diagnostic behavior much more closely than either pure goal-directed or data-directed search (Elstein et al., 1978). The diagnostic process is one of looking at available symptoms, allowing them to suggest higher-level hypotheses, and then setting out to prove or disprove those hypotheses, all the while recognizing hypotheses that might be suggested by symptoms appearing in the verification process. Pauker and Szolovits (1977) have noted that a physician will go to great lengths to explain data inconsistent with a partially verified hypothesis before abandoning it. This type of behavior is not altogether inconsistent with the strategy we have employed, albeit for a different reason. The combination of a partially verified hypothesis and data inconsistent with it may be enough to boost an assertion that would explain the inconsistent data "above" an alternative hypothesis on the agenda. Oddly enough, some of this behavior seems to be a natural consequence of the control structure we have employed.

24.3.1 Generalizing

There is no reason to suppose that the representation and control mechanisms used in WHEEZE could not be used to advantage in other diagnostic production systems. A system similar to EMYCIN (Chapter 15), having both knowledge acquisition and explanation capabilities, could certainly be based on frames and agenda-based control. It also seems likely that an analogue of the EMYCIN rule compiler could be developed to take portions of an assertion network and produce efficient LISP code that would perform identically to the agenda-based control scheme operating on the assertion network.

A second class of extensions that becomes possible with a frame-based system is the addition of other kinds of knowledge not essential to the diagnostic process. For example, in the development of GUIDON (Chapter 26) Clancey noted that a substantial proportion of the domain knowledge had been compiled out of the rules used by most high-performance systems. Within our framework there is no reason why this information could not be added while still maintaining high performance. Such additional information might also be useful for enhanced explanation of system behavior.

24.3.2 Some Outstanding Questions

In the discussion above, claims were made about the perspicuity of the frame representation and about the flexibility of the agenda-based control mechanism. Of course, the acid test would be to see how well domain experts could adapt to the representation and to see whether or not they would become facile at tailoring control flow.

A second question that we pondered is this: how would WHEEZE be different if we had started with a basic frame system and the agenda-based control mechanism and worked with an expert to help build up the system from scratch? It is entirely possible that the backward-chaining production system paradigm had a significant effect on the vocabulary and knowledge that make up both PUFF and CENTAUR. In other words, the medium may have influenced the "message."

To a large extent, we have only paraphrased PUFF's rules in a different representational medium. This paraphrase may not be the most natural way to do diagnosis in the new architecture. Unfortunately, we do not have sufficient expertise in pulmonary function diagnosis to consider radical reformulations of the domain knowledge. For this reason, it would be interesting to see a new diagnostic system developed using the basic architecture we have proposed.