

PROTOTYPES AND PRODUCTION RULES: AN APPROACH TO KNOWLEDGE REPRESENTATION FOR HYPOTHESIS FORMATION

Jan S. Aikins
Heuristic Programming Project, Department of Computer Science
Stanford University, Stanford, CA 94305

A system called CENTAUR has been implemented to interpret data derived from pulmonary function tests using a knowledge representation that combines the advantages of both production rules and frames. The system uses a hypothesis-directed approach to problem solving, in which hypotheses are suggested by the initial data, further information is acquired, and then more specific hypotheses are selected. The hypotheses are represented as **PROTOTYPES**, frame-like data structures each of which characterizes some pulmonary disease. The prototypes guide the invocation of the production rules and focus the search for new information. Some of the advantages afforded by representing knowledge as both prototypes and rules are also presented.

1 Introduction

Much of Artificial Intelligence research has focused on determining the appropriate knowledge representations to use in order to achieve high performance from knowledge-based systems. The principal Artificial Intelligence theme being explored in this present research^{*} is that there are many advantages to a system that uses both frame-like structures and production rules to perform problem-solving tasks in knowledge-intensive domains.

In order to test this theme, a knowledge representation was designed using a combination of frames and production rules. The frames are called Prototypes because they represent stereotypical situations which can be used as a basis for comparison to the actual situation given by the data.^{**} The domain chosen was that of pulmonary physiology. The task was to interpret a set of pulmonary function test results, producing a set of interpretation statements and a diagnosis of pulmonary disease in the patient. A system called CENTAUR has been written to perform this task using prototypes that characterize the typical features of each pulmonary disease. Each feature is called a Component of the prototype. Associated with each component are production rules used to infer a value for the component. These production rules are a form of procedural attachment with a constrained, stylized syntax that makes them easier to examine than general procedures. This constrained syntax leads to other advantages, such as ease of acquisition and modifiability as discussed in [2]. The prototypes focus the search for new information by guiding the invocation of the production rules and eliciting the most relevant information from the user. These prototypes are linked together in a network in which the links specify the relationships between the prototypes.

This research developed out of the MYCIN project [5], which uses a knowledge base of production rules to perform infectious disease consultations. Initially, a MYCIN-like

production rule system called PUFF [4] was written to perform pulmonary function test interpretations. Problems with the production rule formalism in PUFF and similar rule-based systems, such as a need to focus the search for new information and the desire to represent characteristic patterns of disease, motivated the creation of a prototype-directed system.

2 The CENTAUR System

CENTAUR produces an interpretation of data and a pulmonary diagnosis based on a set of pulmonary function test results. The inputs to the system are the pulmonary function test results and a set of patient data including the patient's name, sex, age, and a referral diagnosis. The output consists of both a set of interpretation statements that serve to explain or comment on the pulmonary function test results and a final diagnosis of pulmonary disease in the patient.

CENTAUR uses a hypothesis-directed approach to problem solving where the possible hypotheses are represented by the prototypes. The goal of the system is to confirm some of the prototypes as matching the data in an actual case. The final set of confirmed prototypes is the system's solution for classifying the data in that case. The prototypes represent the various pulmonary diseases, their degrees and subtypes, with the result that the set of confirmed prototypes represents the diagnosis of pulmonary disease in the patient.

In the example below, the prototype representing a pulmonary function consultation itself, the PULMONARY-DISEASE prototype, has been selected as the initial Current Prototype, the system's best hypothesis about how to classify the data in the case. The initial data is requested and the user responses (preceded by a double asterisk **) are recorded. Data entered in the system suggests or "triggers" disease prototypes. The triggered prototypes are placed on a Hypothesis List and ordered according to how closely they match the data. The system attempts to fill in values for the components of the current prototype, which may cause rules to be invoked; or, if no rules are associated with the component, the system will ask the user for the value. When all of the prototype components have values, a decision is made by the system as to whether the given data values match those expected for the prototype. Another prototype is then selected as the Current Prototype, and the process repeats. The system moves through the prototype network confirming or disproving indicated disease prototypes. Matching data and prototypes continues until each piece of data has been accounted for by some confirmed prototype or until the system has concluded that it cannot account for any more of the data.

^{*} This work was supported by the Advanced Research Projects Agency under contract MDA 903-77-C-0322. Computer facilities were provided by the SUMEX-AIM facility at Stanford University under NIH grant RR-00785. The author is sponsored by the Xerox Corporation under the direction of the Xerox Palo Alto Research Center.

^{**} The term prototype has been given the same meaning by other researchers, for example in KRL [1], a prototype is a special kind of unit representing a hypothetical individual that is the typical member of a class.

3 CENTAUR Example

The following is an example of an interpretation of a set of pulmonary function test results for one pulmonary patient. Comments are in italics.

CENTAUR 14-Jan-79 13:54:07

-----PATIENT-7-----

1) Patient's identifying number: ** 7446

2) referral diagnosis: ** ASTHMA

3) FEV1/FVC ratio: ** 40

[Trigger for OAD and CM 900]

(Prototype OAD is triggered by the value 40 for the FEV1/FVC ratio. The Certainty Measure (CM) indicates on a numerical scale the degree of certainty with which the prototype is indicated by the data.)

4) TLC observed/predicted: ** 139

5) FVC/FVC-predicted: ** 81

[Trigger for NORMAL and CM 500]

(The questioning continues and other prototypes are triggered by the data values.)

MoreSpecific Prototypes chosen: NORMAL OAD

(Although there are five possible, more specific disease prototypes for the PULMONARY-DISEASE prototype, only the two that were triggered by the initial data are selected as possibilities to pursue. These prototypes are filled in with the data values that are already known in the case.)

!Surprise Value! 261 for RV in NORMAL, CM: 700

!Surprise Value! 139 for TLC in NORMAL, CM: 400

(Any data values that are not consistent with the values expected for that disease prototype are noted as Surprise Values, and the CM for that prototype is lowered. Two of the data values that are not consistent with the NORMAL pulmonary function prototype are shown here.)

Hypothesis List: (OAD 990) (NORMAL -699)

(The Hypothesis List of triggered prototypes is then ordered according to the CM of the prototypes and a new Current Prototype, OAD, is chosen.)

Components of OAD to trace: F25 D-RV/TLC

(In order to instantiate the OAD prototype, two more components must have values. These are then asked of the user if there are no rules that can be used to deduce their values. The OAD prototype is confirmed as matching the data in this case. Control information associated with the prototype specifies that the Degree of OAD should be determined next, followed by the Subtype of OAD.)

Confirmed: ASTHMA SEVERE-OAD OAD

(Eventually SEVERE-OAD and ASTHMA are also confirmed. Data values that can be accounted for by one of the confirmed prototypes are marked. If there are data values remaining that cannot be accounted for by the confirmed prototypes, the system will attempt to determine if there are multiple diseases in the patient. Refinement Rules associated with the confirmed prototypes are executed to further refine the diagnosis and conclusions which are then printed.)

Conclusions:

Smoking probably exacerbates the severity of the patient's airway obstruction.

Good response to bronchodilators is consistent with an asthmatic condition.

...

Pulmonary Function Diagnosis:

Severe Obstructive Airways Disease.

Asthmatic type.

4 Prototypes and Components

Following frame terminology, each prototype contains SLOTS of information associated with it. One of these is the slot COMPONENTS that lists the substantive characteristics of the prototype. Each component may, in turn, have slots of information associated with it. In the OAD prototype in Figure 4.1, there are components for many of the pulmonary function tests that are useful in characterizing a patient with OAD. For example, the TOTAL LUNG CAPACITY of a patient with OAD is typically higher than that of a person with normal pulmonary function. Thus there is a component, TOTAL LUNG CAPACITY, with a range of PLAUSIBLE VALUES that are characteristic of a person with OAD.

Some control information is represented explicitly in slots associated with the prototype. These slots contain a set of one or more clauses that express some action to be taken by the system in order to instantiate the prototype (CONTROL slot), upon confirmation of the prototype (IF-CONFIRMED slot), in the event that a prototype is disproved (IF-DISPROVED slot), and in a clean-up stage after the system processing has been completed (ACTION slot).

PROTOTYPE	Obstructive Airways Disease
GENERAL SLOTS	Author: Aikins
--Bookkeeping Information	Date: 27-OCT-78
	Source: Dr. Fallat
--Pointers to other prototypes	Pointers:
(link prototype)	(degree MILD-OAD) ...
--English phrases	(subtype ASTHMA) ...
	Hypothesis: "There is an interpretation of OAD."
CONTROL SLOTS	If-Confirmed:
Control	Deduce degree of OAD
If-Confirmed	Deduce subtype of OAD
If-Disproved	Action:
Action	Deduce OAD findings
	Print OAD findings
COMPONENTS	TOTAL LUNG CAPACITY
Plausible Values	Plausible Values: >100
Default Value	Importance: 4
Possible Error Values	
Rules	REVERSIBILITY
Importance of value to this prototype	Rules: 19, 21, 22, 25
	Importance: 1

FIGURE.4.1 A sample prototype with possible slots on the left and values for OAD on the right.

4.1 Production Rules

The CENTAUR knowledge base also includes sets of production rules. Many of the production rules are classified as INFERENCE RULES, rules used to infer information in the domain. They refer to values for components in their premise clauses and make conclusions about values of components in their action clauses. An example of one of the Inference Rules is given in Figure 4.2. The RULES slot associated with a component contains a list of all Inference Rules that make a conclusion about that component. These may be applied when a value is needed for the component.

If: 1) A: The mmf/mmf-pred is less than 20, and
 B: The fvc/fvc-pred is greater than 80, or
 2) A: The mmf/mmf-pred is less than 15, and
 B: The fvc/fvc-pred is less than 80
 Then: 1) There is evidence that the degree
 of OAD is severe, and
 2) One of the OAD findings is:
 Low mid-expiratory flow is consistent
 with severe airway obstruction.

FIGURE.4.2 A Sample Production Rule-English Version

5 Control Structure

The control information used by CENTAUR is contained either in slots associated with the individual prototypes or in a simple interpreter. Control strategies specific to an individual prototype are represented in slots associated with that prototype, with more general system control being expressed in the interpreter.

The control structure can be broken into three stages: a hypothesis-formation stage in which data values are acquired and an attempt is made to match prototypes to data, resulting in a list of confirmed prototypes; a refinement stage in which a set of REFINEMENT RULES are applied to the list of confirmed prototypes to "debug" this list and further refine the recommendations that will be made; and a final "clean-up" stage in which, for example, findings associated with the prototype are printed.

6 Advantages of the Prototype-Directed Approach

The use of this approach for the pulmonary function interpretation task, as compared to the purely rule-based approach used in PUFF, results in two categories of advantages: those dealing with the knowledge base representation itself and those dealing with the system's reasoning and performance. Advantages in knowledge representation occur partly because some knowledge previously represented in rules is now represented more clearly in prototypes. For example, the prototypes explicitly represent control information formerly represented in the PUFF inference rules. In the PUFF system, there are rules whose purpose it is to guide computation by controlling the invocation of other rules. This feature can be very confusing to the medical experts since they do not know which rules are those intended to represent descriptive medical expertise and which rules are those serving a necessary computational function. New knowledge is also being represented in prototypes; for example, plausible ranges of values for each of the pulmonary function tests for each disease can be listed, as well as the relative importance of each measurement in a particular disease prototype. Advantages in system reasoning and performance, that is, the questions that are asked and the order in which information is acquired, include the following:

(A) Consultation flow more closely follows physician's reasoning. The process of medical problem solving has been discussed by many researchers (e.g., [3]) and it is widely felt that a sequence of suggesting hypotheses, acquiring further information, and then revising the hypotheses, as is used in CENTAUR, is, in fact, the problem-solving process used by most physicians. Thus CENTAUR offers increased conceptual clarity, in that the user can understand what the program is doing, and this factor leads to other advantages, for example, the system can offer the user a more intelligible explanation of its performance during the consultation.

(B) Questions are asked in a reasonable order. In a rule-based system such as PUFF, questions are asked of the user as rules are invoked containing clauses referring to information not yet known. The expert can control the order in which the questions are asked only by writing rules to enforce some order. As has been discussed, this procedure results in a potentially confusing rule base where some rules guide computation. In the prototype-directed system, the expert can specify the order in which information is acquired for each prototype in the control slot.

(C) Only relevant questions are asked. Another advantage of CENTAUR over PUFF is that only those hypotheses suggested by the initial data are explored. For example, if the Total Lung Capacity (TLC) for the patient is 70, then CENTAUR would begin exploring the possibility of Restrictive Lung Disease (RLD) because a low TLC would trigger the RLD prototype. (A low TLC is consistent with a hypothesis of RLD; a high TLC is consistent with OAD.) In the PUFF program, the first disease tried is always OAD, so the PUFF program would begin asking questions dealing with OAD. These questions would seem irrelevant considering the data, and, indeed, if there were no data to indicate OAD, such questions would not be asked by CENTAUR.

7 Summary

CENTAUR was designed in response to problems that occurred while using a purely rule-based system. By changing the knowledge representation to include prototypes as well as production rules, new knowledge was represented. Further, knowledge that had been represented rather awkwardly in rules was represented more clearly in prototypes. The production rules were retained as a stylized form of procedural attachment that could be easily examined or modified. By altering the control structure so that a best-fit matching process of prototypes to data produced a current best hypothesis to guide further search, a more focused consultation resulted which more closely followed the way physicians reason. Control knowledge was explicitly labeled and made prototype-specific so that control of the consultation was adapted to the current best hypothesis. In summary, the prototype-directed system achieved better reasoning and performance than the rule-based system. In addition, although representing knowledge in production rules alone did not seem adequate for this task, the ability to represent knowledge in prototypes as well did provide the needed flexibility.

References

- [1] D. G. Bobrow and T. Winograd, An Overview of KRL, a Knowledge Representation Language. *Cognitive Science* 1(1).
- [2] R. Davis and J. King, An Overview of Production Systems. In E. W. Elcock and D. Michie (Eds.), *Machine Intelligence 8*. New York: Wiley & Sons, 1977. Pp. 300-332.
- [3] A. S. Elstein, L. S. Shulman, and S. A. Sprafka, *Medical Problem Solving--An Analysis of Clinical Reasoning*. Harvard University Press, Cambridge, Mass., 1978.
- [4] J. C. Kunz, R. J. Fallat, et al., *A Physiological Rule Based System for Interpreting Pulmonary Function Test Results*. HPP-78-19, C. S. Dept., Stanford University, Dec. 1978.
- [5] E. H. Shortliffe, *Computer-Based Medical Consultations: MYCIN*. New York: American-Elsevier, 1976.