Anthony Andrade's Profile of CASNET

WPI Computer Science Department
Spring Semester 2000

CS 538 Expert Systems
Profile form by Prof. Dave Brown

THE PROFILE QUESTIONS

GENERAL

Domain:
e.g. Chemistry, Electronics, Computers, Medicine, ... (plus specific sub-area).

Medicine(glaucoma)

Main General Function:
e.g. Diagnosis, Prediction, Planning, Data Interpretation, ...

Diagnosis of disease state and recommendation of therapy

System Name:
The CASNET (Causal-Associational NETwork) Project

Dates:
1971 - 1978

Researchers:
A. Safir, C. A. Kulikowski, S.M. Weiss

Location:
Rutger's University, Mount Sinai School of Medicine

Language:
Consolation Program: FORTRAN
Modeling Program: SNOBOL

Machine:
PDP-10 running either TOPs-10 or TENEX operating systems

Brief Summary:

CASNET was built to test a new model for diseases (specifically glaucoma). This model was built by experts to reflect the way that they think when diagnosing a disease. The model was also special in that it represented the disease through time. Using this

model CASNET would determine the progression of glaucoma in the patient and recommend treatment. The system would also be able to track the effects of the treatment and revise it based on its effects on the glaucoma.

Related Systems:

a) None
b) EXPERT, Prospector
c) None

CATEGORY TWO

Characterization of Givens:
What is the information:

Given: Observations of the doctor/ Test Results
Built-in: Model of glaucoma

Characterization of Output:

Output is a diagnosis and suggestions for treatment
Sample Output:

****
DIAGNOSIS AND THERAPY
****

*** RIGHT EYE: ***
[1] PRESENT DIAGNOSTIC STATUS:

(PRIMARY) ANGLE CLOSURE GLAUCOMA. POSSIBLE CILIARY BLOCK MECHANISM. PREVIOUS ACUTE ANGLE CLOSURE ATTACK. DISPARITY BETWEEN THE ANGLES OF THE TWO EYES SUGGESTS POSSIBILITY OF SECONDARY MECHANISM IN ONE EYE. TRAUMA IS A POSSIBLE CAUSE OF THE OPEN ANGLE IN THE FELLOW EYE. POSSIBLE SECONDARY CAUSES OF THE NARROW ANGLE IN THIS EYE: SWOLLEN OR DISLOCATED LENS, TUMOR OF THE CILIARY BODY, INFLAMMATION, OR A CILIARY BLOCK MECHANISM.

[2] TREATMENT RECOMMENDATIONS:

PERIPHERAL IRIDECTOMY IS INDICATED. INSTITUTE WEAK MIOTIC THERAPY AS NEEDED TO NORMALIZE TENSIONS. (ADD DIAMOX IF NECESSARY)

Characterization of Data:

Is the data reliable?
The data consists of current popular beliefs of the mechanisms of glaucoma. Their may be different interpretations by different experts, all of which are provided in the output. Since the model is based on the way the experts think, it can only be as reliable as the group of experts.

Is the data complete?
The model is not assumed to be complete (contradictions in diagnosis may occur).

Generic Tasks:

Which Generic Tasks it obviously includes, explicitly or implicitly?
Classification(diagnosis)
Hypothesis matching (uncertainty handling methods).
Heuristic classification because it does not do data abstraction.

Theoretical Commitment:
Does the system have any theoretical underpinning? Is it claiming to show that some theory of its type of problem-solving is correct? Is the method used claimed to work for other similar domains?

(see next answer)

Reality:

Is there any psychological validity to the method used the structure of the knowledge, the control mechanisms? That is, is it a system that is merely a simulation of result or is it in any way a simulation of method at a cognitive level?

Yes, the system is designed to emulate expert thinking. The idea behind using this model was to move away from (theory-based) probabilistic reasoning and towards a system of reasoning that experts could relate to and feel comfortable using.

CATEGORY THREE

Completeness:

Has the system been fully implemented?

Yes.

Use:

Has the system been used with real users from outside the original

development situation?

The system was a prototype, but the information gained from building it went into designing EXPERT.

Has the system been used with real users in the user's own working environment?

No

Performance:

Are there any performance measures available?

Yes

How was the system evaluated? How did it fare?

The evaluations during a 1976 opthamalogical conference were subjective, but over 77% of those questioned rated the system as very competent or better. There was also a study in 1978 where CASNET was tested on twelve complex cases; eleven of which it correctly diagnosed. The incorrect case had region-specific influences that the experts who built the model may have not had much exposure to.

CATEGORY FOUR

Phases:

Is the system organized into distinct phases of different activity? Distinct subtasks? What are they?

The consultation system has two main subtasks: interpret the observations/results and gather new information. CASNET takes the observations about the patient and matches them against the model to loosely determine the state of the patient. The system then tries to refine its interpretation by asking for more observations (e.g. lab tests). These are both part of diagnosis.

Subfunctions:

Despite the fact that the system has a single main function (e.g., diagnosis) does it use other types of problem-solving as part of the system? (see the answers to the last question for clues as to whether this answer might be yes). What kind(s)?

e.g. Diagnostic systems sometimes select therapy too; this can be done by synthesis, or by selection (which is often a kind of

classification that is, can this disease situation be classified as one that this drug can help?)

The system also recommends a treatment. Once the disease state and observations(e.g. allergies) are known it is a selection problem.

Use of Simulation or Analysis:

Does the system use a numerical simulation or analysis, either done by itself or by some package, during its operation?

CASNET uses weights between nodes to do calculations, but these weights are considered symbolic.

System/Control Implementation Architecture:

i.e. The overall architecture
e.g. meta-rules + rules, blackboard + knowledge sources, production rules, active agents, activation nets ...

The model of CASNET is a causal-associational network. The network has four layers: observations, p-states, disease states, and therapies. The system starts with observations and then moves through p-states toward a diagnosis and therapy guided by the evidence given to it. The second half of the model consists of normative rules to control weight calculations.

CATEGORY FIVE

Characterization of Structure Knowledge:

i.e. Is it grouped into types? What types? What are the types used for? Do they correspond to phases of the system?

e.g. component knowledge, chemical knowledge, functional knowledge, causal knowledge, ...

The network has four layers: observations, p-states, disease states, and therapies.

Characterization of Process Knowledge:

i.e. For active knowledge, how would you characterize the effect of that knowledge?

e.g. <Partial Situation Description> --> <Classification>.

<patient data> → <p-state> → <disease category> → <therapy>
Patient data is analyzed to suggest which p-states are present. The p-states will characterize a disease with a certain severity. Therapy will be selected corresponding to that disease and severity level.

Deep or Surface:

i.e. Is the system using deep knowledge or is it only using surface knowledge? How would you characterize it? Are there levels of representation or reasoning?

e.g. If there is a qualitative simulation then it could be argued that it is using deep knowledge.

The many layers of the model indicate that the system contains deep knowledge of the mechanisms of glaucoma.

CATEGORY SIX

Search Space:

What space or spaces does the system search through? Is the search explicit? Are the states represented explicitly? How? What do the states represent? Complete alternative solutions, Solution refinements, Plans, ...? How big is the space?

CASNET searches through the four categories of the model, discussed previously. The states are represented explicitly. Observations represent the symptoms and test results of the patient. The p-states represent internal problems that indicate specific diseases. The disease states indicate a disease and a particular level of severity of that disease. The therapy states represent a class of treatments. There are alternative solutions in the form of alternate opinions on the diagnosis and treatments. The state space consists of more than 100 states, 400 tests, 75 classification, and 200 diagnostic and treatment statements.

Space Traversal:

How is the space traversed? What does it mean, in terms of the problem, to move from one part of the space to another? Sub-problem decomposition, gradual refinement, down pre- determined hierarchy, instantiation, ...

Weights are calculated to provide levels of confidence of a state. These weights are combined with the weights of the arcs (which indicate the confidence that one state

causes another) to produce the confidence of the occurrence of the next level. The system will produce a hypothesis and recommendation based on the confidence levels of each state. If the confidence of the hypothesis is not acceptable it will ask for more information and adjust the weights of the nodes based on these levels.

Search Control Strategy:

Does the system use a strategy that is expressible in terms of the problem? or does it appear just to be an AI "technique" that happens to fit?

The system uses a causal-associational network to reason. It is not just an AI "technique" that happens to fit. It was chosen because it is a rational way to reason and it simulates a thought process that is believed to occur in humans.

Standard Search Strategies:

Does the system explicity use Generate & Test or Means-Ends Analysis?
(N.B. be careful, almost everything is a form of G&T!)

A diagnosis is generated from the observations. It is tested against an acceptable confidence level. If it fails the process is repeated with more observations.

Search Control Characterization:

e.g. depth first, best first, breadth first, knowledge-based, random, ...

Casnet's search could be classified as a best first, since it evaluates the nodes and then focuses on refining those that will give it the most benefit.

Subproblems:

Is evaluation of partial solutions possible?

(that is, can the system know when it is on the right track?)

Partial Solutions can exist in the sense that a disease state and treatment can be found without the desired level of confidence.

Are the subproblems independent? totally? partially?

No. Finding a node in each level of the mode requires the weights of the previous level. Recommending a treatment depends on diagnosis being completed.

Search Control Representation:

Where is the search control knowledge? How is it expressed? Explicitly?

The control knowledge is expressed in the model as normative rules for calculating weights of arcs in the model.

Search Control Strength:

Is it based on a very domain independent and knowledge-free method (a "weak" method) or is it very domain dependent and knowledge-full (a "strong" method).

Control knowledge is domain independent and knowledge-free, a weak control method.

CATEGORY SEVEN

Failure Method:

When part of the system fails, or reaches an incorrect conclusion, how does the system attempt recovery? What type of knowledge does it use? is it global or local? Is there a significant vari- ation from the normal flow of control? If it uses back-tracking, what kind does it use?

Contradictions are noted and described to the user. Alternative explanations are also offered for topics that are not fully understood yet.

Uncertainty:

What is it that is uncertain?

- **a piece of data?**
- **a piece of knowledge?**
- **a solution to a subproblem?**

A piece of data may be uncertain. This is represented via a three-value true system, explained in the next answer. Alternative explanations on disputed topics are also provided.

Management of Uncertainty:

Does the system use probabilities, scoring values, a fixed range of certainty values, ...? What does a value mean? What is the method of combination? Is it a global method, or does it vary? (i.e., a local method). Are there any apparent problems?

CASNET uses a three-value true system (true, false, and unknown). Scoring values range from –1 to 1. A value is a numeric equivalent to a symbolic value (i.e. a 1 means

definitely confirmed). The method of combination is to initially set all observation values to unknown. Then when new values are added, the confidence weights only change if the new observation brings evidence that has a more significant impact than the original weight. This process is used to propagate the data through the entire model. One possible drawback is that the confidence factor is only indicative of its most influential

Management of Time:

Is there any time-dependent data? How does it affect the problem-solving?

Yes. Time-dependent data allows monitoring of treatment and also to predict future developments of the patient.

CATEGORY EIGHT

Knowledge Representation Method:

e.g. rules, procedures, tabular, semantic nets, logic, ...

Knowledge is represented by the model. The causal and associational links comprise the theoretical knowledge, while the normative component contains the practical rules for applying data to the knowledge.

Knowledge Representation Generality:

Is there a special language for this system or does it use a general method?
Is the method provided by some Expert System building tool?

The model may be modified using an interactive editing program, so the language it was written in was hidden from users.

Knowledge Structuring:

Is the system based on a hierarchy, a network, ...? Does this structure correspond to the domain in some way? Does it correspond to the problem-solving being carried out?

The knowledge structure is a causal-associational network. It represents the relationships between the different elements within the domain. The network is set up so that it is traversed in the same way an expert would think about a problem.

CATEGORY NINE

Alternative Representations:

Does the system use alternative representations for the same piece of knowledge in order to allow for alternative solution methods?

Yes. It allows multiple explanations for the same diagnosis. It also allows alternate therapy recommendations based on certain patient observations (allergies, age).

Alternative Solution Methods:

Does the system use alternative methods to reach the same solution(s)?

A p-state state can be confirmed with equal strength by multiple observations. The same applies to a disease state being confirmed via multiple p-states. The system may also offer alternate therapies (where curing the patient would be considered the solution). The same applies to offering multiple tests to gather the needed data.

Optimization:

Does the system produce the best answer? Sometimes? Always?

No. The system can produce contradictions because the model is assumed incomplete.

Multiple Results:

Does the system produce more than one result? If it obtains several does it try to order them by evaluating them? Does it try to combine them? If so, how?

The system gives one diagnosis and treatment class, which could contain multiple explanations and could be refined with new observations. If a treatment class is viewed as multiple results, then the system orders them by its preference based on the facts it knows about the patient (age, allergies, etc).

CATEGORY TEN

Interaction:

i.e. Is there anything interesting about the way the system interacts with the user? (for both input and output).

e.g. menus, pictures, diagrams, picking, color, sound, ...

The system uses text input and output and interacts in a question-answer manner.

Data collection:

Does the system require all its data (that is, details of this particular problem) before execution or does it allow/require incremental addition? If incremental, does it vary the type/number/ order of data gathering actions (e.g., questions of a data-base or user) depending on the problem being solved?

The system allows incremental addition of data into it. It tries to limit the number of the questions and the cost of the steps needed to obtain the answers.

Data format:

i.e. In what form is the data given?
e.g. patient records, LISP, Natural Language, ...

The data is given as answers to the system's questions using as small of a range of numbers as possible. Or yes/no/unknown answers.

Acquisition:

Does the system have any way of acquiring knowledge from the expert user? Does it guide the user? Does it have a way of validating the knowledge?

A nationwide group of experts (ONET) constantly updated the system to refine the model. The system guides the user with a series of questions. It asks for further tests or more questions to try to improve the confidence it has in a hypothesis.

Learning:

Does the system learn from its own performance

No.

Explanation:

Does the system have the ability to explain where its result came

from? Is this obtained from a trace of the goals and sub goals formed during system execution?

Yes. The traces the path it took through the network and uses this as the basis for the explanation.

CATEGORY ELEVEN

Strengths:

What do you think are the strong points of the system?

CASNET's model is its major strength:

- It is designed to simulate the way experts think.
- It is intuitive and consequently can offer intuitive explanations.
- It includes temporal aspects of disease diagnosis and therapy.
- It is separate from the data meaning it can be updated without affecting the rest of the system
- It separates the descriptive and normative components so that one may be updated without affecting the other.

ONET continuously refined the model so the system was always up to date

Weaknesses:

What do you think are the weak points of the system?

The system did not have a strong theoretical basis, so proving the validity decisions of the system would be more difficult.

The system is only as good as ONET. This was seen when it failed to correctly diagnose a glaucoma that occurred only in a high percentage of Japanese men. It can also be seen in the fact that contradictions can occur because the model is incomplete.

Other:

Any other comments that would help to characterize this system.