THE ROLE OF EXPERIENCE IN DEVELOPMENT OF EXPERTISE

Janet L. Kolodner

School of Information and Computer Science
Georgia Institute of Technology
Atlanta, GA 30332

INTRODUCTION

Perhaps the most distinguishing feature of an expert is that when given a novel problem to solve in his or her domain of expertise, the expert can solve the problem easily. Novices, on the other hand, are good at dealing with typical problems or "classic" cases, but not novel problems. In people, the evolution from novice to expert happens as a result of being able to introspect and examine the knowledge used in solving problems. That introspection and examination allows people to learn from experience. A human expert can interpret a new case in terms of something (either a previous case or generalized knowledge) he is already familiar with. This implies that as an expert is having new experiences, he is evaluating and understanding them in terms of past experiences. In the process, he is integrating the new experience into his memory so that it too will be accessible to use in understanding a later case.

AN EXAMPLE

The following example [7] illustrates experiential learning in the psychiatric domain:

> Dr. X sees a patient who seems to show classic signs of major depression. She is 38 and complains of depression. In the last month, she has been unhappy, suffering from insomnia and crying. She reports poor concentration and diminished interest in activities, and that she has been depressed since childhood when her father deserted the family. She has previously been diagnosed as depressive, and was treated in a mental hospital with antidepressants. She was sickly as a child, has had a drinking problem, and has had a number of physical illnesses which doctors have not been able to find causes for.
>
> Seeing that she has been treated previously for depression, that her chief complaint is depression, and that she has insomnia, poor concentration, and diminished interest, Dr. X concludes that this patient is suffering from Major Depression, Recurrent, without Melancholia. He treats her with antidepressants. The antidepressants seem to work, but the woman keeps coming back complaining of additional major physical disorders. Dr. X begins to think that there may be some other problem which he had not accounted for.
>
> Taking a further history, he finds out more about the medical problems she has had. They have been numerous, and doctors have been unable to find organic reasons for them. He realizes that the large number of medical problems is important to consider. Going through the diagnostic process again using that symptom as the predominant clinical feature, he realizes that he should have diagnosed her for Somatization Disorder in addition to the diagnosis of Depression.

As a result of this case, Dr. X should learn the following:

1. it is important to take medical history into account in choosing predominant clinical features,
2. depression can camouflage somatization disorder, and
3. a patient who is highly depressed but who complains about medical problems may be suffering from somatization disorder in addition to depression

Using the first fact, he should be able to refine his rules for choosing predominant clinical features. This case should help him conclude that medical history may be a more important clinical feature than he had previously realized. In addition, the next time he finds that a treatment has failed, he may be biased towards finding out facts about the patient's medical history that he hadn't known. Using the second fact, he should be skeptical of diagnoses of depression coming from other doctors, and will want to find out more about the medical history of a new patient before taking a previous diagnosis seriously. The third fact gives the relationship between depression and somatization disorder, which could be helpful in diagnosing and treating later cases.

Furthermore, this case should enable the doctor to hypothesize that there is often more than one symptom that is primary, and that a patient's chief complaint may not be the most important symptom to look at. This experience should also help him learn that current problems must be separated from long-standing problems and both must be taken into account. Later cases should enable him to

*This research has been supported in part by NIH-BRSG Grant #5 S07 RR 07024-16 and in part by NSF Grant #IST-8116892.

recognize and separate current from long-standing problems, to confirm these two hypotheses, and to learn specifically which long-standing problems and which chief complaints are likely to be significant.

THE SHRINK PROJECT

We are investigating the relationship between experience and the development of expertise in the domain of psychiatry. When complete, our program, SHRINK, will analyze new psychiatric cases based on previous cases it has seen. It will integrate new cases into its memory as it is processing them, and will build up its expertise based on generalizations it has made concerning the similarities between cases it has seen (similarity-based learning, [2] [3]). Its expertise will also be heightened through analysis of failures in diagnosis and treatment (failure-based learning, [6] [4]).

A recently-published psychiatric manual [1] provides a procedure for diagnosis. Knowledge about symptoms and knowledge about diseases are integrated into the process. Related disorders are organized in decision trees. Predominant symptoms (called primary indices) suggest entry points into those trees. The book also specifies sets of necessary, supporting, additional, and exclusion criteria for each currently-recognized psychiatric disorder. We are starting with the processes specified in that book as a model of a novice diagnostician, and giving our program experiences which allow it to reorganize that knowledge "episodically" and learn. The following is SHRINK making a diagnosis.

```
--- Patient Background and Information ---

The patient is Ms. 163, a female, 25 years old.
She has been exhibiting the following symptoms:
   attempted suicide by CARBON-MONOXIDE
   SIGNIFICANT DECREASE in state S-WEIGHT
   SIGNIFICANT DECREASE in state S-MENTAL-FOCUS
   SIGNIFICANT DECREASE in drive D-SLEEP
   less than NORMAL on the HAPPINESS scale

**** Begin Diagnostic Impression ****

Deciding predominant clinical symptom:
     attempted suicide by $CARBON-MONOXIDE
Starting differential diagnosis at:
     DD-MOOD-DISTURBANCE
   Possible diagnoses are:  (MAJOR-DEPRESSION)

**** Begin Diagnostic Evaluation ****

     ...
Symptoms confirm MAJOR-DEPRESSION

Ms. 163 suffers from:
   MAJOR-DEPRESSION supported by:
     SINGLE-EPISODE-MAJOR-DEPRESSION supported by:
       DEPRESSIVE-EPISODE
```

Currently, SHRINK knows about Major Depression and diagnoses normal cases of single-episode depression and manic depressive syndrome using the novice procedure. In addition to making diagnoses, the machinery for learning by experience is being put into place. The system indexes cases it has seen, and can pick out the similarities and differences between individual cases. It keeps track of the decisions it makes in doing its diagnoses and records those decisions and the reasons it had for making them. We have also given it rudimentary treatment capabilities. SHRINK is now ready to learn. The first example it will work on is the one above.

MEMORY STRUCTURES

In refining task expertise, two different types of knowledge must be learned or refined — domain knowledge used by the reasoning process, and the reasoning process itself.

In order to incrementally refine diagnostic procedures (i.e., the reasoning process), memory must have an explicit model of the procedures it is employing. We call these memory structures PROCESS MOPs (Memory Organization Packets, [5] [2]). The prime PROCESS MOP is DIAGNOSE. The first step in psychiatric diagnosis involves examining the patient and choosing predominant symptoms. Probable disorders are chosen based on that set of symptoms (differential diagnosis). Each is then evaluated in more detail and unlikely diagnoses are deleted (diagnostic evaluation). Failures in the diagnostic evaluation can suggest additional disorders which must be evaluated.

The diagnosis process is driven by diagnostic knowledge -- knowledge about particular disorders. The structures which record knowledge about particular disorders are called **DIAGNOSTIC MOPs.** The figure below shows some of the information SHRINK has about "Depressive Episodes", one of its diagnostic MOPs. Presence of a "depressive episode" is necessary to diagnose all "major affective disorders", including "major depression" (two more of its DIAGNOSTIC MOPs).

A third type of knowledge memory must have is symptom knowledge -- knowledge of how particular symptoms and their combinations tend to predict particular diagnostic categories. Symptoms suggest entry points into differential diagnosis decision trees, thus suggesting general diagnostic categories a patient may fit into. Attempted suicide, for example, suggests severe mood disturbance, which suggests a possibility of depression.

Each step of the diagnostic process is guided by either disorder or symptom knowledge. After predominant clinical features (major symptoms) are chosen, the knowledge associated with them is used to choose starting points for differential diagnosis. Only those categories implied by the primary symptoms are considered in initial differential diagnosis. Differential diagnosis is guided by knowledge about the relationships between diagnostic categories. Disorder knowledge (in particular, knowledge about the normal symptoms and

exclusion criteria associated with particular disorders) also guides diagnostic evaluation once possible disorders have been established through differential diagnosis.

DEPRESSIVE EPISODE

ESSENTIAL FEATURES:
  at least 1 of:
    (1) dysphoric mood
    (2) pervasive loss of interest or pleasure in usual pastimes and activities.
SUPPORTING FEATURES:
  at least 4 of:
    (1) a significant change in appetite.
    (2) a disturbance in sleep habits.
    (3) psychomotor agitation or retardation.
    (4) decrease in sexual drive.
    (5) loss of energy or fatigue.
    (6) feelings of worthlessness, self-reproach, or excessive or inappropriate guilt
    (7) complaints or evidence of diminished ability to think or concentrate.
    (8) recurrent thoughts of death, suicide, death wishes, or suicide attempt.
  time constraint: symptoms must have been present simultaneously and for at least 2 weeks
EXCLUSION CRITERIA:
 (1) preoccupation with a mood-incongruent delusion or hallucination (=> PSYCHOTIC DISORDER)
 (2) bizarre behavior (poss => PSYCHOTIC DISORDER)
ADDITIONAL CRITERIA:
   negate diagnoses of:
     (1) SCHIZOPHRENIA
     (2) SCHIZOPHRENIFORM DISORDER
     (3) ORGANIC MENTAL DISORDER
     (4) UNCOMPLICATED BEREAVEMENT
TREATMENT:
   choose a combination of
     (1) antidepressant treatment
     (2) hospitalization
     (3) ECS therapy
     (4) analysis

INCREMENTAL LEARNING

In order for failures in procedure to refine previously-held diagnostic rules, memory's organization must be updated with each new experience. For that to happen in the memory structures we have defined, two things must happen when an experience deviates from the expected. First, the deviant path through appropriate PROCESS MOPs must be recorded. Second, deviant features of this case with respect to previous similar cases must be recorded in DIAGNOSTIC MOPs.

There are two reasons for recording these differences. First, they should be recorded so that when a similar deviation occurs later, the original case can be remembered. The two cases can then be examined for similarities, and generalizations based on their similarities can be used to evolve a new diagnostic category. This results in **similarity-driven learning.** When later similar cases occur, the knowledge necessary to deal with them will already be in memory. Second, a deviation signals that additional reasoning must be attempted. Extra effort is then applied to "explain" the deviation. If an explanation is found, diagnostic procedures are updated, and that case is maintained as support of the new procedure. If no explanation is found, the deviation marks a problem that must be resolved later. Later, when a similar problem case is encountered, the marker signals that both the old case and the new one should be examined to see if between them they provide enough evidence to explain the violation. This is the process of **failure-driven learning.** Markers in process MOPs allow incremental process changes. Markers in diagnostic MOPs allow refinements in diagnostic categories.

The Example in More Detail

To see how we aim to get the computer to learn incrementally, we will consider the example above in more detail.

In order to make the major depression diagnosis, the doctor must have the following knowledge:

1. Depression, when reported by a patient, means severe unhappiness.
2. Severe unhappiness is a form of mood disturbance.
3. Mood disturbance (a symptom) suggests that the patient suffers from major depression.

This knowledge allows a differential diagnosis of "major depression" to be made. We also assume that the doctor knows and can recognize symptoms of a depressive episode. He recognizes three supporting symptoms for a depressive episode in addition to mood disturbance -- insomnia, poor concentration, and diminished interest. Because the patient has previously been diagnosed as depressive, he is willing to conclude that she has had a depressive episode even though only three symptoms are present (normally, diagnosis of a depressive episode requires four symptoms). The doctor thus concludes that the patient has Major Depression, Recurrent, without Melancholia. Note that because a rule was stretched, this judgement may be suspect, and that if a failure occurs, this step will have to be examined as a possible reason for the failure. We assume that this doctor knows that a common treatment for major depression is antidepressants, and therefore prescribes that treatment. As a result of the medication, her depressive symptoms go away, the expected successful result of the medication.

In this case, however, there is also an additional response to the medication. The patient complains of a number of physical symptoms. This is where failure-based reasoning comes in. This is unexpected, and must be explained. There are two problems to look at here — the process of noticing violations and the process of explaining them. A violation of an expectation may be a reaction completely different than the expected one, or there may be varying degrees of resemblance. Alter-

natively, as in this case, a violation is recognized even though the expected reaction is present, since there is additional behavior that cannot be accounted for.

The problem of "assigning blame" for expectation violations or failures is a hard problem both for people and for computers. In this case, there are two things to explain. First, causes for the new problem must be found. Second, the doctor must find if the new problem could have been prevented (i.e., what, if anything, failed in the diagnosis or treatment process).

There are three possible causes for the patient's new complaint:

a. the medicine is producing side effects,
b. the patient has developed a new physical ailment, or
c. the patient is imagining her disorders.

Each of these possibilities must be checked out. We assume the first doesn't hold. Considering the next possibility, appropriate diagnostic tests should be performed on the patient. After taking further history, performing a physical examination, and doing screening tests, if no organic reason for the illness is found, then the doctor should be "reminded" [5] [2] of the patient's previous medical history. This should happen because this experience is similar to the previous medical experiences the patient has reported to the doctor -- she has a history of physical illnesses for which doctors have been unable to find causes. This should lead the doctor to take a further history to check the possibility of Somatization Disorder.

Note that there are many types of "reminding" that need to be done in the processing we are suggesting. The reasoner (whether human or computer) should be reminded of previous cases similar to the one it is currently working on so that knowledge gleaned from those cases can be used to reason about the new case. Such reminding is required, for example, in relaxing diagnostic rules. In addition, the reasoner must be reminded of episodes associated with the case it is currently working on, so that it can notice patterns of behavior associated with the current patient. The implication here is that memory must be organized along both of these dimensions.

Reminding must also allow a third type of knowledge to be remembered — symptom knowledge. Recall that when the doctor realizes that the patient is probably imagining her symptoms and that she has a long history of imagining physical disorders, he hypothesizes that she might be suffering from Somatization Disorder. In order for him to make that hypothesis, he must be reminded of the following piece of symptom knowledge:

4. Excessive medical history in females implies a somatization disorder.

Why can the doctor remember this rule now when he didn't remember it during the initial diagnosis? Our claim is that knowledge only becomes available when there is a way to direct memory processes toward it. He did not remember this rule initially because he was not focussing on the medical history of the patient in trying to come up with predominant clinical features. Instead, he was focussing on the then current problem. Because he was focussing entirely on the mood disturbance, only knowledge associated with mood disturbances was accessible for reasoning. Only when attention is directed to the patient's physical disorders and medical history as possible clues to her illness, does the knowledge associated with those symptoms become available. Memory must be organized so that relevant knowledge is available when attention is directed to it, and so that irrevelant knowledge does not get in the way of reasoning.

Using the symptom knowledge listed above, the doctor can now hypothesize that the patient has a Somatization Disorder in addition to depression. He does a differential diagnosis and diagnostic evaluation based on that hypothesis and finds that her past medical history does support the hypothesis.

At this point, the doctor has corrected his initial mistake. In order to learn from it, he must figure out whether and why he made the mistake. In general, the problem of figuring out where in the process a mistake was made is hard. In the domain of diagnosis and treatment, there are four possibilities:

a. the diagnosis was wrong, and therefore the treatment is unsatisfactory,
b. the diagnosis was right, but the treatment was not appropriate to the diagnosis
c. the diagnosis was right, but the treatment didn't work, or
d. the treatment and diagnosis were right, but something new has come up.

In finding the cause of the new complaint, the doctor has already found that an additional diagnosis had to be made, so the original diagnosis could have been wrong (a). On the other hand, he might not have had the necessary information initially to make this diagnosis. Therefore, (d) could also be the case.

To distinguish between the two of these, the doctor must decide whether he had the necessary information in the beginning to make the correct diagnosis. The key to making that decision is figuring out where in the diagnosis there might have been a problem and how it could have been corrected. Diagnosis consists of a number of processes and in general, a reasoning failure may happen during any of them. In this case, however, there is a direct route to finding the initial failure. Having corrected the mistake, the diagnostician has the crucial piece of information that allows him to figure out where in the process he went wrong. He knows that the symptom knowledge

in rule 4 would have been necessary to initially diagnose Somatization Disorder. Furthermore, he knows that symptom knowledge of this sort is part of the initial diagnosis process -- establishing primary indices.

His error, then, was in choosing predominant clinical features. Once he knows where in the process the error occured, he must determine whether he had enough information initially to include Somatization Disorder originally. He did have this knowledge, since the patient had already reported having been sickly and having had a number of illnesses for which no organic causes could be found. He concludes that he should have paid attention to that initially, and marks "medical history" as a patient feature which he should pay more attention to in the future.

In addition to concluding that medical history is an important clinical feature to consider in diagnosis, something more specific is learned — the relationship between somatization disorder and depression. This will be represented in a number of ways in episodic memory. First, there will be markers or indices associated with major depression. One way this case differed from normal major depressive cases is that the patient had a large number of previous medical illnesses. That patient feature will be one of the features which index this individual case among the major depression cases diagnosed and treated. If another case comes up in which the patient is depressive and also has had a number of previous illnesses, then this case, which included an additional diagnosis of somatization disorder, will be remembered. This case will also be indexed as one in which the patient began complaining of other illnesses after treatment for depression seemed successful. If another similar case comes up, the doctor or system can be reminded of this case. That should cause him to wonder whether there is a previous medical history which he had not elicited from the patient, and if the second patient also has somatization disorder. Episodic memory will contain similar markers associated with somatization disorder relating it to depression.

DISCUSSION

There are two major implications of the model we have presented. First, it is set up to handle exceptional cases very nicely. The general case is stored as a diagnostic category. In any but the exceptional cases, the general knowledge associated with the diagnostic category is used for diagnosis and treatment. As exceptions are encountered (e.g., the somatization case above where normal depressive treatment was inadequate), they are indexed off of the diagnostic category. If an explanation of the exception has been made, it is stored along with the exceptional case. When a new case is reminiscent of a previous exceptional case, knowledge about the previous case can be used to deal with the new case. When an exception has been encountered and dealt with successfully a number of times, it evolves into a new diagnostic category with its own specialized diagnostic and treatment rules. The general rules associated with the original diagnostic category do not change, however, unless an exception becomes the generalized case.

Another advantage this approach has is that it deals with both experiential knowledge and facts in the same way. Both are stored in the same structures and organized identically. This means that both are equally accessible and both can be used in reasoning. An implication of this organization is that it is amenable to new information from any source. New treatments and methods of diagnosis discovered through experience and those learned through journal articles or from others can be added to memory in the same way. Hopefully, the same processes used for reorganizing memory due to failure in experience can be used to reorganize memory based on new information acquired elsewhere. Of course, as in people, only through experience will all the implications of such knowledge be learned and added to memory.

SUMMARY

This paper has pointed out how experience aids in developing the expertise necessary for expert reasoning. It has also introduced a computer program based on these ideas. The research and the program are still in a state of infancy. Nevertheless, we see this approach as having a great deal of promise both in terms of implementing expert computer systems and in helping us to understand the cognitive processes underlying expertise.

ACKNOWLEDGEMENTS

Dr. Robert M. Kolodner, a psychiatrist at the Atlanta VA Medical Center, has been the major informant for this project. Keith McGreggor has done all of the programming. Both have been helpful in formulating the ideas put forth in this paper.

REFERENCES

[1] American Psychiatric Association (1980). Diagnostic and Statistical Manual of Mental Disorders (Third Edition). American Psychiatric Association, Washington, D.C.

[2] Kolodner, J. L. (1980). Retrieval and organizational strategies in conceptual memory: a computer model. Ph.D. thesis. Research Report #187. Dept. of Computer Science, Yale University, New Haven, CT.

[3] Lebowitz, M. (1980). Generalization and Memory in an Integrated Understanding System. Ph. D. Thesis. Research Report #186. Department of Computer Science. Yale University, New Haven, CT.

[4] Riesbeck, C. K. (1981). Failure-driven reminding for incremental learning. In Proceedings of IJCAI-81.

[5] Schank, R. C. (1980). Language and memory. Cognitive Science, Vol. 14, No. 3.

[6] Schank, R. C. (1981). Failure-driven memory. Cognition and Brain Theory. Vol. 1.

[7] Spitzer, R. L., Skodol, A. E., Gibbon, M., and Williams, J. B.W. (1981). DSM-III Case Book. American Psychiatric Assoc., Washington, D.C.