MORE: An Intelligent Knowledge Acquisition Tool

Gary Kahn, Steve Nowlan, and John McDermott

Department of Computer Science

Carnegie-Mellon University

Pittsburgh Pennsylvania 15213

Abstract

MORE is a tool that assists in eliciting knowledge from domain experts. Acquired information is added to a domain model of qualitative causal relations that may hold among hypotheses, symptoms, and background conditions. After generating diagnostic rules from the domain model, MORE prompts for additional information that would allow a stronger set of diagnostic rules to be generated, MORE'S primary value lies in its understanding of what kinds of knowledge are likely to be diagnostically significant. By formulating its questions in a way that focuses on such knowledge, it makes the most effective use of the domain experts' time.[1]

I. The MORE system

MORE elicits diagnostically significant knowledge from domain experts; it is similar in spirit to systems like TEIRESIAS [Davis 82] and ETS [Boose 84]. Like these systems, MORE provides a mechanism for interviewing domain experts and makes use of several strategies for facilitating the interview process, MORE differs in that it takes a model-theoretic approach to the acquisition of diagnostic knowledge. It uses a qualitative model of causal relations together with a theory of how causal knowledge can be used to achieve more accurate diagnostic conclusions to guide the interview process. Thus, it can go farther in diagnosing weaknesses in a knowledge base than its predecessors. As these weaknesses are identified, MORE elicits from the domain expert information that can lead to stronger diagnostic conclusions.

MORE evolved from our experience handcrafting MUD, a diagnostic system in the drilling fluids domain [Kahn 85]. In reflecting on the nature of our interactions with the drilling fluids experts, it seemed to us that we had used a set of quite broadly applicable knowledge acquisition strategies, MORE embodies those strategies.

MORE has the capacity to build domain models from a fixed set of qualitative relations that may hold among hypotheses, symptoms, and background conditions. The content of any particular model is provided by a domain expert in response to Its prompts, MORE generates diagnostic rules from the domain model. After a rule is constructed, the user is asked to associate positive- and negative-support values with each rule. MORE'S procedure for constructing and evaluating rules is described elsewhere. [Kahn 84]. The focus of this paper is on the nature of the domain model MORE constructs and on MORE'S knowledge acquisition strategies.

MORE plays three roles while interacting with a user. Its first role is as an information solicitor. As the user enters the names of hypotheses and symptoms, he is asked for additional information that may result in a stronger diagnostic assessment. Among other things, MORE will ask about events that may affect the expectation of the hypothesis occurring, the likelihood of seeing the symptom given that the hypothesis has occurred, and distinguishing characteristics of the symptom that may identify it as having been caused by a particular hypothesis.

Once an initial knowledge base is built up, MORE looks for weaknesses in the rules it has generated. As a result of evaluating the diagnostic strength of these rules, MORE prompts for additional information that would allow a stronger set of diagnostic rules to be generated. For instance, if there is no rule assigning a high positive-support value to a symptom (S) that bears on hypothesis (H), MORE asks if there are distinguishing characteristics of (S) when it is caused by (H). This strategy is called symptom distinction, it is one of eight strategies discussed below.

In it's third role, MORE looks for potential inconsistencies in the way a user has assigned confidence factors to diagnostic rules, determining if the assigned weight is appropriate given the weights of other rules already in the system. Although MORE does not have sufficient information to recommend precise value assignments, it knows enough to formulate expectations regarding the direction in which weights ought to vary across different rules. For instance, if background condition (C) is known to increase the likelihood of observing symptom (S) when hypothesis (H) occurs, then the negative-support of a rule whose evidential focus is (S) and just (S), should be lower than the negative-support of a rule whose evidential focus is (S) and (C). If the user assigns weights that violate this expectation, MORE issues a warning and asks the user if he wishes to change the weights of any of the conflicting rules.

[1] Mott of the work reported here was done while all three authors were at Carnegie-Meilon University. Gary Kahn is now at Carnegie Group Inc.

2. MORE's domain model

As domain experts are interviewed by MORE, a representation of their responses is built into a 'domain model'. Each domain model built by MORE consists of five representational entities: hypotheses, symptoms, conditions, links, and paths. A hypothesis denotes an event whose identification will be the result of a diagnosis. A symptom is any event or state consequent to the occurrence of a hypothesis, and whose observation disposes toward the acceptance of the hypothesis. A condition is an event or state in the environment which is not directly symptomatic of any hypothesis but which can affect the diagnostic significance of some other event. Links are used to join entities in the model, including other links. A path is a special type of link that joins a hypothesis to a symptom. States or events represented as symptoms may also be hypotheses. Figure 2-1 provides a schematic representation of the key representational objects and relations used by MORE.

Five kinds of conditions can be represented in a MORE model: frequency-conditions, tests, test-conditions, symptom-conditions, and symptom-attributes. Frequency-conditions are used to represent anything that can affect the a *priori* expectation of a given hypothesis. Such conditions are assumed to be independent of observing any particular symptom. Thus, in a MORE model, they are linked directly to an affected hypothesis.

Tests represent procedures or devices used to determinine the occurrence of a symptom; test-conditions represent events or states that bear on the accurate use of a procedure, device, or observation. Since the expected accuracy of a detection procedure may affect the evidential significance of a symptom with respect to any of its explanatory hypotheses, tests are linked directly to their corresponding symptoms. As test-conditions are conditions which affect the confidence in the results or accuracy of a given test with respect to a particular symptom, they are attached to the link joining a test and symptom.

Symptom-attributes are specific characteristics of a symptom that tend to make it more or less likely to be caused by a particular hypothesis. An example would be the contiguity of datacheck reports for a disk, if the datachecks were on contiguous sectors this would indicate a radial scratch more strongly than if the datachecks were randomly distributed. Symptom-attributes provide a way of refining the description of a symptom into one of a number of subclasses, each subclass providing greater discrimination among causes of the symptom. Since these conditions refine a symptoms description with respect to the kind of thing which could cause it, they are attached to the path connecting the hypothesis and symptom.

Symptom-conditions represent states or events which affect the likelihood of a symptom occurring if the hypothesis has. External events which could, for example, mask or preclude the realization of a symptom even if the hypothesis occurs are represented as symptom-conditions. Since these conditions affect the causal link between a hypothesis and symptom, they are attached to the path connecting the hypothesis and symptom.

3. Strategies for improving diagnostic performance

MORE'S suggestions on how to augment the existing knowledge base such that stronger and more accurate diagnostic conclusions can be obtained result from its use of eight different

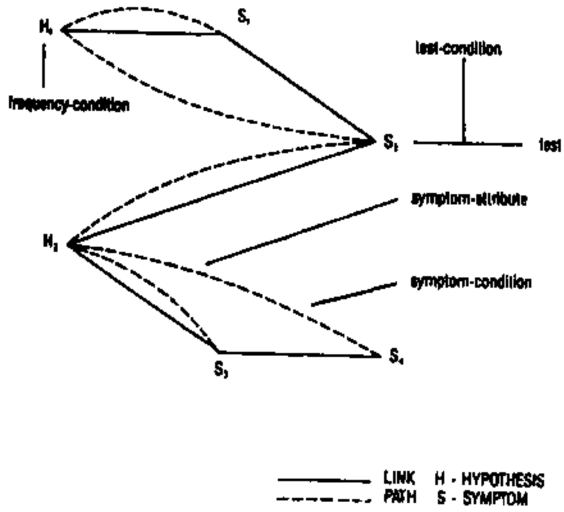

Figure 2-1: MORE objects and relations

strategies:

- differentiation
- frequency conditionalization
- symptom distinction
- symptom conditionalization
- path division
- path differentiation
- test differentiation
- test conditionalization

The differentiation strategy leads MORE to actively seek symptoms that provide leverage in distinguishing among diagnosable events, MORE scans its current domain model looking for pairs of hypotheses for which there is no differentiating symptom. When it finds such a pair, the user is ask for a symptom that will differentiate the pair of hypotheses. With respect to MORES model a symptom (S) is said to differentiate one hypothesis (H1) from another (H2) when there is a path from H1 to S, and no path from H2 to S. Since increased differentiation in the knowledge base also results from incorporating symptoms which are explainable by a set of causes different (at least in part) from those underlying previously reported symptoms, MORE also scans for triples of hypotheses (H1,H2,H3) for which there is no symptom which differentiates H1 from both H2 and H3. The identification of such symptoms makes it more unlikely that a set of symptoms could be explained erroneously by two co-occurring hypotheses.

For instance, in the MUD domain, both an influx of water and an insufficient use of emulsifier can have the same effects on measurable mud properties. However, an increase in mud volume is usually associated with the former. While this effect can also result from a hydrocarbon influx, other shifts in mud properties distinguish hydrocarbon from water influxes. Thus, the knowledge base can be further differentiated by adding the fact that an increase in volume is a confirming observation with respect to a water influx.

Path differentiation is another means of finding symptoms with greater diagnostic significance. Under this strategy, MORE asks if a symptom, which may result from one of several causes, does so via (at least partially) non-overlapping causal pathways. Intermediary events on non-overlapping portions of these pathways are expected to have greater diagnostic significance than symptomatic events on shared pathways. Thus, if there is a link to S from both H1 and H2, MORE asks about the existence of more proximal symptomatic events which lie on the the causal path from H1 to S, but not on the path from H2 to S. Such symptoms would be associated with stronger positive-support values as there are fewer hypotheses which could explain their occurrence.

In the MUD domain, for example, an increase in plastic viscosity in an oil mud can result from either shale or water contamination. These effects, however, do not result in entirely the same way. Shale contamination causes an increase in plastic viscosity by increasing the percentage of solids in the mud system; water causes an increase by its behavior in a partially emulsified solution. The mud engineer can determine which of these mechanisms accounts for increased plastic viscosity through the use of additional tests. These tests measure the amount of unemulsified water and the solids content of the mud. Positive results on these tests provide stronger confirmation of the respective causes than does the shared symptom of increased plastic viscosity.

Path division similarly requires eliciting a symptomatic event that lies on a causal path from the diagnosable event to an already reported symptom. In this case, however, the new symptom is selected such that it is more expected, given the cause, than the former more distal symptom. Path division is attempted when no rule with a symptom (S) immediately linked to a hypothesis (H) has a high negative-support value. When this is recognized, MORE prompts for a symptomatic event (S2) that lies on the causal pathway from H to S, such that S2 is more proximally caused by H and causes S. As such, the failure to observe the new proximal symptom will be of greater disconfirmatory value *ceteris paribus* than failing to observe symptoms later in the causal chain.

An increase in bentonite, for example, can be considered an intermediate step between shale collapsing into the bore hole and a change in viscosity. Thus, as expected, it was found that the failure to observe an increase in viscosity is less disconfirmatory with respect to shale contamination than the failure to observe a significant increase in free bentonite through the use of a methylene blue test.

Symptom conditionalization requires seeking out events that effect the likelihood that a symptom will occur, given that the hypothesis has. Negative-support values typically vary with the expectation that a diagnosable event will indeed give rise to a particular symptom. This expectation can be low if, for instance, the appearance of a symptom requires the cooccurence of an unlikely background condition.

For example, some viscosity effects normally associated with salt contamination of a water based drilling fluid will appear only if the fluid has not been pretreated with surfactant thinners. If there has been a pretreatment of this kind, the failure of viscosity symptoms to appear cannot count as evidence against the hypothesis of salt contamination. However, if one knows that the system has not been pretreated in this way, then the disconfirmatory significance of failing to observe these viscosity symptoms is much greater than it would be otherwise.

Symptom distinction seeks to further distinguish a symptom so it may more closely be identified with a particular cause. Under this strategy MORE asks about symptom-attributes which are capable of dividing a symptom into finer classes, with each class more likely to be associated with one particular cause (or hypothesis) than any other. As a symptom is distinquished by characteristic attributes, the number of explanatory hypotheses goes down, and the positive support contributed by the observation of such a symptom goes up. Similarly, if a symptom is always characterized by a certain attribute when it is caused by a particular hypothesis (H), then the observation of the symptom without these characteristic attributes should lead to a stronger disposition to reject H. For instance, both an influx of water and an increase in low specific gravity solids can cause a decrease in density. However, if density has decreased rapidly, it is more likely to have been due to an influx of water.

Both symptom-conditions and symptom-attributes are represented in MORES domain model as conditions attached to the path joining the hypothesis and symptom. This permits similar attributes to have different effects with respect to different hypotheses.

In test differentiation MORE seeks to determine if there are procedures or observational instruments that can detect the existence of the symptom with greater accuracy. If there are, the expectation is that reference to them will increase the evidential significance of the symptoms on which they bear. Tests, as represented in MORE'S model, are attached by a link directly to the symptom whose evidential significance may be affected. The significance of changes in pH level, for instance, differ slightly depending on whether pH is measured by litmus paper or the more accurate pH meter.

Similarly, in pursuing test conditionalization, MORE inquires about conditions which can affect the confidence in the results or accuracy of a test. These conditions are represented as test-conditions in the model and are attached to the test whose accuracy they modify.

Finally, MORE will, under certain conditions, engage in frequency-conditionalization. With this strategy, MORE looks for conditions that affect the expected likelihood of a hypothesis' occurrence. These are represented as frequency-conditions in the domain model, and are attached directly to the hypothesis they modify. Frequency-conditions are mapped into rules which will enhance or diminish the measure of belief provided by a set of symptoms, depending on whether these conditions increase or diminish the expectation of the hypothesized cause. For instance, in the MUD domain, an increase in viscosity often results from drilling through one of a number of contaminants, some of which may be expected, others unexpected, in the location being drilled. Thus, one would like the evidential significance of a symptom, such as an increase in viscosity, to be dependent on local knowledge about the likelihood of encountering various contaminants.

4. A little concreteness

MORE seeks diagnostically significant information by prompting its user with questions such as:

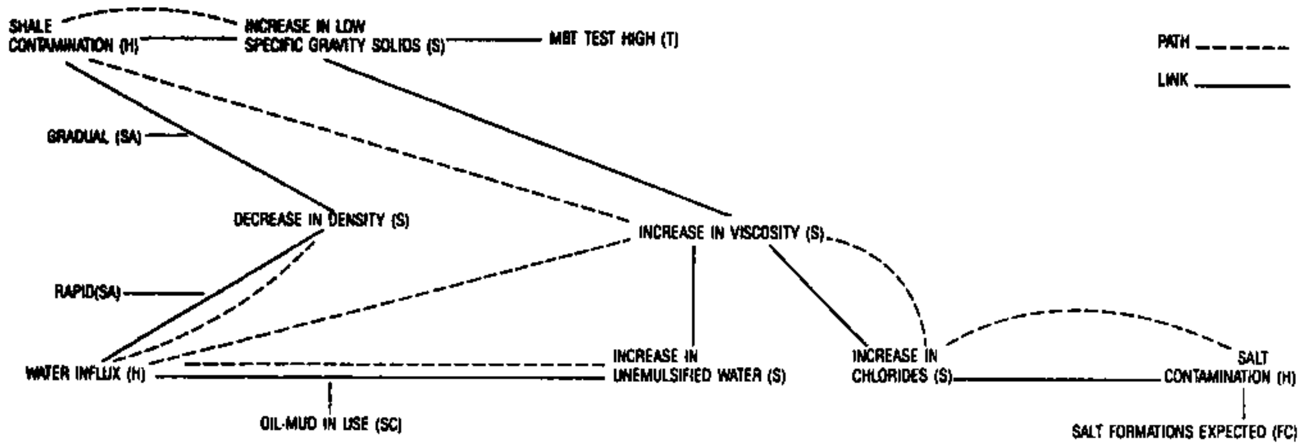

Figure 4-1: A MORE model

1. Are there any conditions under which the problem referred to by SHALE-CONTAMINATION would be more or less likely to occur?
2. Are there any conditions that affect the accuracy (or confidence) of observing INCREASEINVISCOSITY?
3. Are there any distinguishing features of DECREASE-IN-DENSITY which would make it more or less likely to be caused by WATER-INFLUX?
4. Can you provide a symptom associated with WATER-INFLUX that cannot be explained by SHALE-CONTAMINATION?

The above questions are driven respectively by strategies of frequency conditionalization, test conditionalization, symptom distinction, and differentiation, MORE uses a number of heuristic rules in deciding when it is appropriate to pursue one strategy or another [Kahn 84].

As diagnostic knowledge becomes available, MORE maps it into its underlying domain model. For example, in discussion above, several hypotheses, symptoms, and conditions relevant to the diagnosis of drilling fluid problems were mentioned. These included three hypotheses (shale contamination, water influx, and salt contamination), five symptoms (an increase in low specific gravity solids, an increase in viscosity, a decrease in density, an increase in unemuslified water, and an increase in chlorides), two symptom attributes (a gradual and a rapid decrease of density), a symptom condition (the use of an oil mud), a frequency condition (the expection of salt formations), and finally, a test (a high MBT reading). Within the current implementation of MORE, each of these is represented as a separate data structure linked within a network of causal relations as illustrated in figure 4-1.

5. Concluding remarks

MORE assists in knowledge-base construction using the eight strategies described above. In the course of its development, MORE has been applied to parts of the drilling fluids domain as well as to sample diagnostic problems provided by a physician. Our next step is to use MORE to develop a number of knowledge-baaed consultation systems in a wide variety of domains, MORE is currently being used to build systems to diagnose computer disk faults, computer network problems, and circuit board manufacturing problems. These efforts should give a good indication of the power of MORE'S strategies.

References

[Boose 84] Boose, J.
Personal construct theory and the transfer of human expertise.
In *Proceedings of the National Conference on Artificial Intelligence.* Austin, Texas, 1984.

[Davis 82] Davis, R. and D. Lenat.
Knowledge-Based Systems in Artificial Intelligence.
McGraw-Hill, 1982.

[Kahn 84] Kahn, G., S. Nowlan, and J. McDermott.
A foundation for knowledge acquisition.
In *Proceedings of IEEE Workshop on Principles of Knowledge-based Systems.* Denver, Colorado, 1984.

[Kahn 85] Kahn, G. and J. McDermott.
MUD, a drilling fluids consultant.
Technical Report, Carnegie-Mellon University, Department of Computer Science, 1985.