STRUCTURE OF MEDICAL KNOWLEDGE AND DEEP REASONING(Research supp

by National Library of Medicine Career Development Award K04 LM00083, National Science Foundation grant MCS 8305032, NIH Grant R01 LM 04298 from the National Library of Medicine, a NIH Training Grant T15 LM 07023 from the National Library of Me

B. Chandrasekaran, Jon Sticklen, and Jack W. Smith, Jr.,
The Ohio State University, Columbus, Ohio 43210 USA.

Extended Abstract: In this talk we will provide a historical/conceptual account of the evolution of expert systems in general, and medical diagnostic expert systems in particular, from the viewpoint of the degree to which such systems embody capabilities for functional or "causal" reasoning. We distinguish between deep models in the sense of scientific first principles and deep cognitive models where the problem solver has a qualitative symbolic representation of the system or device that accounts qualtitatively for how the system ''works.'' We analyze diagnostic reasoning as an information processing task, and identify the generic types of knowledge (and reasoning) that are needed for the task to be performed adequately. If these types of knowledge are available, then we show how an integrated collection of generic problem solvers can produce a diagnostic conclusion. The need for deep or causal models arises when some or all of the knowledge of the types identified above are missing in the problem solver. We provide a typology of different knowledge structures and reasoning processes that play a role in qualitative or functional reasoning. We indicate where the work of Kuipers, de Kleer and Brown, Davis, Forbus, Bylander,

Sembugamoorthy and Chandrasekaran fit in this typology and what types of information each of them can produce. We elaborate on functional representations as deep cognitive models for some aspects of causal reasoning in medicine.

Role of Causal Reasoning in Diagnosis: Typically, a diagnostic problem starts with the observation of some behavior which is recognized as deviation from the expected or desirable, i.e., a malfunction behavior is observed. The problem solver at this stage needs to generate some hypotheses about the cause of the malfunction: typically these are in terms of changes in the structure of the device from the specifications. At this stage in many domains (e.g., the medical domain) a number of low-cost broad spectrum testing (e.g., physical examination, a battery of blood tests) may be undertaken without any specific hypotheses in mind, or the initial malfunction may be used to invoke one or more specific malfunction hypotheses. Most often these hypotheses are invoked by using what one might call ''precompiled'' pieces of knowledge that relate behavioral observations to one or more hypotheses. This initial hypothesis generation task can be more or less complex, and more or less controlled depending upon the domain, and the knowledge the problem solver has. Whatever the particular method, they all involve going from behavioral observation (test values, signs and symptoms, etc.) to a number of hypotheses, possibly ranked.

At this stage typically a small number of the more plausible hypotheses are considered the differentials. In a compiled system, knowledge may be explicitly available for each hypothesis in the differential about which further tests may be useful for confirmation or rejection of that hypothesis, and in that case by comparing this knowledge for the different hypotheses in the differential, the problem solver can generate tests that have the potential for the greatest discrimination between the hypotheses. If, however, this knowledge is not directly available to the problem solver, but the structure of the device is known, then the following reasoning can be very useful. Assume the structure change corresponding to each of the malfunction hypotheses in the differential list, and reason about what behavior will follow. Often, for a number of reasons, one would like to do this qualitatively. The basic work in this area was initiated by de Kleer, and he and Brown at Xerox PARC, Forbus, and Kuipers are among those who have continued this line of work. Reasoning from a given structure to its behavior is required not only in diagnostic reasoning; it is also useful in design, and in planning where the problem solver will need to project the behavior of a design or plan to check conformity to specifications.

Architecture of Causal Reasoning: Causal reasoning about devices or physical systems involves multiple types of knowledge structures and reasoning mechanisms. Two broad types of approaches can be distinguished. In one, causal reasoning is viewed mainly as an ability to reason at different levels of detail: the work of Weiss and Kulikowski, Patil and Pople come to mind. Any hierarchies in this line of work have as organizing principle different levels of detail. In the other strand of work, causal reasoning is viewed as reasoning from structure of a device to its behavior, from behavior to its function, and from all this to diagnostic conclusions. In this approach, the hierarchical organization of the device or system naturally results in an ability to move into more or less levels of detail. For purposes of our current discussion, the following stages can be recognized in this approach to causal reasoning:

1. Given a representation of the behavior of the components of a device or system, and a representation of the structure of the device, i.e. the interconnection of the components, the ability to generate the behavioral description of the device as a whole is an important part of causal reasoning. In simple devices or systems this stage will generate enough information to understand the device. But in general, this technique is useful for producing various fragments of behavior for ranges of values of components. Often these fragments may need to be further organized to explicitly represent the hierarchical structure of the device and also to capture the teleology of the device as in 2 below.

It is to be noted that, as a rule, in addition to behavioral descriptions of components, substantial amounts of domain knowledge general common sense (''naive physics'') knowledge may be needed for this reasoning. Forbus, Kuipers, de Kleer and Brown, and Bylander and Chandrasekaran have presented accounts of how a qualitative account of behavior can be obtained given a structural description of objects and their connectivity. Stanfill has looked at some aspects of common sense spatial reasoning about simple machines.

2. Given the ability to generate behavioral sequences for

various assumptions about the components, the agent can often put together an account of the functions of the device, and its relationship to its structure. In simple cases, the behavior that we talked about in 1 above can be thought of as the function, but in general, functional specifications involve teleology, i. e., an account of how device goals relate to the behaviors of the device. Also, behavior may often need to be abstracted to a level higher than that at which the component is specified. For example, in an electronic circuit, the behavior of the components such as a transistor and a resistor may be in terms of voltages and currents, while a device containing them may be described as an amplifier or oscillator. This abstraction process often involves a hierarchical organization of the relations between function and structure.

Sembugamoorthy and Chandrasekaran have discussed the nature of such a functional representation, and have proposed that it capture in some sense an agent's understanding of how the device functions. In general how an agent constructs a functional account from the structure and behavioral specifications of the components is an interesting theoretical question. de Kleer has provided some examples of this process.

3. While the stages so far help in understanding how the device works, these structures will need to be used in specific ways to help in specific problem solving tasks. The most commonly studied task in this connection is diagnostic reasoning. Often, one can generate diagnostic possibilities (malfunction modes), and test data that will help in determining the presence of these, from one's understanding of how the device works. The paper by Sembugamoorthy and Chandrasekaran outlines how their functional representation can be manipulated by device-independent processes to produce diagnostic knowledge of this type. Sticklen, Chandrasekaran and Smith have work in progress that applies these notions to the medical domain.
