Artificial Intelligence in Medicine 2 (1990) 129–134
Elsevier


# Making deepness explicit

John WASHBROOK and Elpida KERAVNOU

*Department of Computer Science, University College London, Gower Street, London, WC1E 6BT UK*

Received April 1990

**Abstract.** The concept of *deepness* is a useful, if poorly defined, concept. In spite of the development over a number of years of several medical expert systems with high levels of performance, these systems have failed in that they have not been accepted by the medical community. The introduction of so-called **deep systems**, where deepness is often taken to be synonymous with embodying causality, was an approach to resolving this problem. Two systems, CASNET and NEOMYCIN, are compared and it is argued that although CASNET is deeper in that its reasoning is causal, NEOMYCIN is in fact more acceptable in that its explanations and dialogue are closer to those of an expert. The intuitive meaning of deepness is discussed, and a working definition is developed which is not based exclusively upon causality, but which includes the explicit representation of strategic and factual knowledge.

**Keywords.** Medical expert systems, Deepness, Explanation.

## 1. Introduction

### *1.1 The failure of medical expert systems*

There are a number of good medical expert systems in the fields of diagnosis (e.g. NEOMYCIN, INTERNIST, MDX), treatment (MYCIN, CASNET), and patient management (ABEL). Whilst these systems have achieved a high degree of performance, nevertheless it can be argued that none of them have been successful in that their use in daily practice is very limited – they have not been accepted by practitioners.

There are doubtless many and various reasons why these systems are not in everyday use. There are practical considerations such as: the requirement for special computing equipment, poor response times, the lack of access to a medical database and the consequent need to re-enter patient data for every consultation. In addition there may be psychological reasons such as: inflexible user interfaces, consumer resistance to new technology and practices.

However, the two reasons which we consider the most generally relevant and focus on are the **dialogue structure** and the **explanation structure**. By dialogue structure we mean the dynamics of the user-system interaction: its nature, e.g. mixed initiative with the user able to volunteer data and focussing guidance at any time during the consultation; the ordering of the questions; the number, relevancy and comprehensibility of the questions; and whether access to external information (such as databases) is possible. Explanation structure refers to: the way that an explanation is presented; the level and depth of the explanation; and the model of the user (if any) which is used in tailoring the explanation to the particular user.

The dialogue structure and explanation structure are of great relevance as they relate directly to the responsibility which the practitioner has for the well-being of the patient. In order to discharge this responsibility the practitioner must be able to validate any advice offered by the system; this will usually require an understanding of the system's reasoning, which in turn leads to the need for the system to be able to explain its reasoning in a meaningful way. But the majority of expert systems fail to do this; they do not provide acceptable explanations.

### *1.2 Deep systems*

Technology has provided deep systems, with the implicit assumption that they are an improvement on what came before. So, what is a deep system? This question has no clear-cut answer. Firstly, let us consider what a deep system is *not*. A deep system is not shallow – where by a shallow system is meant one in which knowledge is represented as empirical associations, typically if-then rules [7]. Whilst shallow systems may perform well in the majority of cases, they cannot offer an explanation of their reasoning other than by reference to their rules [4, 5]. They cannot reason from first principles, and so cannot descend to a deeper level either to justify the rules or when faced with an exception. This lack of a causal reasoning ability leads us to what is probably the most widely-held definition of deepness in the medical field: that of embodying causality [16]. Other descriptions of the properties of deep systems [18] include temporal reasoning (which is related to causality) [15], qualitative reasoning, reasoning from the first principles, and reasoning from structure and function [13, 4, 5, 11]. We shall argue that deepness is more related to explicitness than causality and there exist systems which are, in this sense, deep (e.g. the Oxford System of Medicine [6, 17]) although such claims might not be made for them – possibly because they are not overtly causal systems. In short, within the medical AI community causality seems to be sufficiently well accepted as embodying deepness that it needs to be challenged.

### *1.3 Deepness as causality*

Philosophers have been arguing about the nature of cause and effect for centuries, and continue to do so. Nevertheless, systems developers manage to build systems embodying causality, and we appeal to the 'commonsense' intuition of causality which most of us seem to share with the developers of those systems. But we claim that causality is by no means synonymous with deepness. To focus our discussion we argue in the next section that CASNET is more 'causal' than NEOMYCIN, but that NEOMYCIN is more 'acceptable' than CASNET and hence must be deeper.

## 2. A comparison of two systems

In this discussion we are concerned with the explanation and dialogue structures of two systems, both of which can be claimed to be deep, without regard to the quantity or quality of any medical domain knowledge. Later we seek a definition of deepness which corresponds better with our notion of acceptability in terms of the explanation structure and dialogue structure.

### *2.1 CASNET – a system embodying causality*

CASNET (Causal ASsociational NETwork) is a diagnostic and treatment system [19]. It has a rich differentiation of knowledge types (planes) into diseases, causes, observations and therapies. The causal plane is the key to diagnosis. It consists of dysfunctional states. The progress of a disease is modelled by causal paths connecting states.

Initial observations, either volunteered by the user or obtained via a protocol, are associated (abductively) with nodes in the causal plane. The causal paths lead to the generation of complex hypotheses, which suggest (deductively) further observations. This cycle of observation-gathering and hypothesis-generation is continued until no further useful information can be gleaned. The confirmed causal states are now related by table-lookup to the disease categories, which in turn point to therapies. Treatment for a specific patient is mediated by the observations, which provide indications/contra-indications. The disease taxonomy is not used during the diagnostic phase, in contrast to NEOMYCIN (see below).

CASNET's reasoning is encoded procedurally in Fortran subroutines. As a consequence CASNET is unable to explain *during* a consultation what it is doing, or why it is pursuing a particular line. However, during the consultation a case-specific picture is created, and at the end of the consultation this can be used to show how progress was made to the solution, and which possibilities were explored.

### *2.2 NEOMYCIN – a system with factual and strategic knowledge*

NEOMYCIN also has a rich set of knowledge types [2]. In contrast with CASNET, where the disease taxonomy is not used during the diagnostic phase, in NEOMYCIN the disease taxonomy is the key structure and reasoning is primarily based upon it. As the disease taxonomy is large a means is needed to narrow down the search and this is provided by **triggers** – cheap findings which can be used to home in on a cluster of diseases. At each node in the disease taxonomy there are Mycin-like (object) rules which use symptoms to refine the disease category. Also described at each node are sequences of symptoms and their duration. Common cases may be diagnosed solely by this structure.

In addition to the disease taxonomy with its top-to-bottom refinement, NEOMYCIN has a causal net (like CASNET's) which provides support knowledge. But in NEOMYCIN reasoning is *backwards* in time, from observed effects to their causes. Thus, common cases are dealt with by the rules (heuristics), whilst exceptions (rare/difficult cases) and explanations of observed symptoms use the causal net to explicate the underlying causal chain. The causal network provides an alternative route for diagnosis, making NEOMYCIN more flexible than CASNET. Like an expert, NEOMYCIN can deal quickly and efficiently with common cases by using heuristics, and handle exceptions by going back to look at causes. This choice of strategies is controlled by a *task taxonomy*, which is represented declaratively, and thus more explicitly, than procedural code. The refinement of a task is controlled by (meta-)rules. Examples of tasks and subtasks are:

- Collect initial data
- Establish hypothesis space
  - Explore
  - Generate and refine
- Process hard data.

As this strategic knowledge is explicit it can be used in explanations [1, 8].

### *2.3 Comparison of CASNET with NEOMYCIN*

The brief descriptions of these two systems (summarised in *Table 1*) shows that CASNET is more 'causal' than NEOMYCIN. CASNET's diagnosis is always based upon causality. It can explain its conclusions, but it cannot explain its reasoning strategy as this is coded in Fortran. For simple cases NEOMYCIN makes its diagnosis in a purely heuristic way, but even so it is more 'acceptable' as it is able to explain its diagnostic strategy and give a rationale for each and every question which it asked (as well as being able to consult its causal network if required).

| Reasoning | |
|---|---|
| CASNET | NEOMYCIN |
| reasoning driven by the causal network | reasoning rule-driven via the etiological taxonomy with causality as an alternative if needed |
| disease taxonomy passive | disease taxonomy active — trigger associations provide focus |
| reasoning knowledge in Fortran routines (procedural) | reasoning knowledge made abstract and explicit (declarative) |
| reasons forwards in time | reasons backwards in time |
| **Acceptability** | |
| CASNET | NEOMYCIN |
| has a mixed initiative dialogue | has a mixed initiative dialogue |
| cannot indicate *what* it is doing, or *why*, during a consultation but *can* explain its conclusions: • why a state was confirmed or denied • why a causal pathway was selected • how conflicting evidence was resolved | has a diagnostic plan which registers • task instantiations • justifications and covers reasoning *during* a consultation, including a rationale for every question. |

Table 1. Comparison of CASNET with NEOMYCIN.

### *2.4 Comparison with other systems*

In [9] six systems are compared in a similar way to that above, with results as shown in *Table 2*. In this table the higher a system the more causal (acceptable) it is. From this it can be seen that CASNET is rated highly for causality but only moderately for acceptability, and conversely for NEOMYCIN.

| Causal | Acceptable |
|---|---|
| CASNET, ABEL | NEOMYCIN |
| NEOMYCIN | ABEL, MDX |
| INTERNIST-1 | CASNET, INTERNIST-1 |
| MDX, MYCIN | MYCIN |

Table 2. Comparison of six systems.

## 3. Deepness as explicitness

Shallow is a pejorative term, whilst deep is evocative – who would prefer a shallow system to a deep one, regardless of the system's function? When we say that people are deep it is their thoughtfulness, the profundity of their insights, which are being referred to. They are able to strip a problem of irrelevant detail, to grasp its fundamentals, and to work rationally towards a solution based on sound principles. When we listen to their explanations we feel that we have gained some new insights, and have confidence in the conclusions which have been arrived at. Presumably these are the qualities which the user expects of a deep system.

Experts are characterized by having a profound and wide knowledge of their domains, and penetrating intellects. When an expert is described as deep it is their skill in bringing these faculties to bear which is being referred to. An expert in diagnosis is able to handle *common* cases efficiently, can reason from first principles in *exceptional* cases, and may be able to explain this train of thought. Thus, when a system is referred to as deep, a user's intuition is that the system will have these desirable properties. But how should an implementor provide deepness? Without a clear working definition, the temptation is to provide solutions which are (relatively) easy or interesting to implement, rather than to provide whatever it is that will address the user's needs. For example, causality is an appealing candidate as it lends itself to representation by directed graphs, a powerful, well-understood technology. In order to avoid this trap we need a working definition of deepness.

### *3.1 Working definitions of deepness*

Klein and Finin [12] give a *relative* definition of deepness: A model $M'$ is deeper than a model $M$ if it represents knowledge, or is able to infer knowledge, which is implicit in $M$. That is, there is more explicit knowledge in $M'$ than in $M$. By this definition, NEOMYCIN is deeper than CASNET: NEOMYCIN's reasoning knowledge is explicit, whilst CASNET's is shrouded in Fortran.

A relative definition allows a comparison of systems, but for a system builder an absolute definition may be more useful. We propose an absolute working definition:

A model is *deep* if

(1) *factual* knowledge is represented explicitly,

(2) *strategic* knowledge is represented explicitly, and

(3) the factual and strategic knowledge cover *exceptions.*

In addition, the system should be *non-monotonic*, that is it should be able to retract conclusions in the light of new facts, for example when a default assumption is proven invalid, and to modify its reasoning plan (retract strategies) dynamically [10].

| Causal | Acceptable | Explicit |
|---|---|---|
| CASNET, ABEL | NEOMYCIN | NEOMYCIN |
| NEOMYCIN | ABEL, MDX | ABEL |
| INTERNIST-1 | CASNET, INTERNIST-1 | CASNET, MDX |
| | | INTERNIST-1 |
| MDX, MYCIN | MYCIN | MYCIN |

Table 3.

This definition proposes necessary conditions for a system to be considered deep. It seems unlikely that a sufficient set of conditions could be proposed. Deepness is a multi-dimensional concept [3], so that a simple ranking of systems on a linear scale of deepness will not always be possible (nor is it with the Klein and Finin definition). However, for the systems considered earlier, using our definition we rank them in order of decreasing deepness:

NEOMYCIN; ABEL; CASNET, MDX; INTERNIST-1

as in *Table 3*, from which we see that the ranking using our definition of deepness corresponds much more closely with acceptability than does causality.

## 4. Summary

Medical expert systems based on shallow knowledge have failed to be accepted. Systems embodying deep knowledge are needed. Causality, by itself, does not give deepness. Rather, a deep system must represent the domain strategic and factual knowledge explicitly; only then can it explain and justify its reasoning and questioning, both at intermediate points during a consultation as well as at the end.

## References

[1] W.J. Clancey, The advantages of abstract control knowledge in expert system design, *Proc. AAAI-83* (1983) 74–78.

[2] W.J. Clancey and R. Letsinger, Neomycin: Reconfiguring a rule-based expert system for application to teaching, *Proc. IJCAI- 81* (1981) 829–836.

[3] A.G. Cohn, Approaches to qualitative reasoning, *Artificial Intelligence Rev.* 3 (1989) 177–232.

[4] R. Davis, Expert Systems: where are we and where are we going?, *AI Mag.* (Winter 1983).

[5] R. Davis, Reasoning from first principle in electronic troubleshooting, *Internat. J. Man-Machine Studies* 19 (1983) 403–423.

[6] J. Fox, A.J. Glowinski, M. O'Neil and D.A. Clark, Decision making as a logical process, *Proc. Expert Systems '88* (Cambridge University Press, UK, 1988).

[7] P.E. Hart, Direction for AI in the eighties, *SIGART Newsletter* 79 (January 1982).

[8] D.W. Hasling, W.J. Clancey and G. Rennels, Strategic explanations for a diagnostic consultation system, *Internat. J. Man-Machine Studies* 20 (1984) 3–19.

[9] E.T. Keravnou and J. Washbrook, Deep and shallow models in medical expert systems, *Artificial Intelligence in Med.* 1 (1989) 11–28.

[10] E.T. Keravnou and J. Washbrook, What is a deep expert system?, *Knowledge Engineering Rev.* 4 (1989) 205–233.

[11] J. de Kleer and J.S. Brown, A qualitative physics based on confluences, *Artificial Intelligence* 24 (1984) 7–83.

[12] D. Klein and T. Finin, What's in a deep model: a characterization of knowledge depth in intelligent safety systems, *Proc. IJCAI-87* (1987) 559–562.

[13] B. Kuipers, Commonsense reasoning about causality: deriving behaviour from structure, *Artificial Intelligence* 24 (1984) 169–203.

[14] B. Kuipers, Qualitative simulation as causal explanation, *IEEE Trans. Systems Man Cybernetics* SMC-17 (1987) 432–444.

[15] W.J. Long, Reasoning about state from causation and time in a medical domain, *Proc. AAAI-83* (1983) 251–254.

[16] P.L. Miller and P.R. Fisher, Causal models for artificial intelligence, in: P.L. Miller (ed.), *Selected topics in Medical Artificial Intelligence* (Computers and Medicine Series, Springer-Verlag, Berlin, 1988),

[17] M. O'Neil, A. Glowinski and J. Fox, A symbolic theory of decision-making applied to several medical tasks, *AIME 89* (Springer-Verlag, Berlin, 1989) 62–71.

[18] C.J. Price and M. Lee, Deep knowledge tutorial and bibliography, *Alvey Report IKBS3/26/048* (1988).

[19] S.M. Weiss, C.A. Kulikowski, S. Amarel and A. Safir, A model-based method for computer-aided medical decision-making, *Artificial Intelligence* 11 (1978) 145–172.