DIALOG: A MODEL OF DIAGNOSTIC LOGIC FOR INTERNAL MEDICINE

Harry E. Pople, Jr., Ph.D.
Graduate School of Business
University of Pittsburgh
Pittsburgh, Pennsylvania

Jack D. Myers, M.D.
School of Medicine
University of Pittsburgh
Pittsburgh, Pennsylvania

Randolph A. Miller
School of Medicine
University of Pittsburgh
Pittsburgh, Pennsylvania

Abstract

A system for computer assisted medical diagnosis has been developed, which incorporates an innovative model of diagnostic logic. A supporting medical data base has also been assembled, now comprising approximately fifty percent of the major diseases ofinternalmedicine. Using weighted associations between disease entitles and their manifestations, and employing a powerful attention focusing heuristic, the system has demonstrated competence In dealing with difficult clinical problems involving multiple diagnoses.

1. Introduction

1.1. Internal Medicine as a Task Environment for Artificial Intelligence

At first glance, the task environment of internalmedicine might appear to be eminently suitable for attack by means of the stste-space methods (1) of artificial intelligence. A medical problem may be characterised In terms of an initial state (disease), a goal state (health), and a collection of available operators (drugs, surgical procedures, diets, etc.) by which one state may be transformed into another. The main difficulty with this view of the task is that ordinarily, evidence available at the outset provides only a partial description of the initial state of a medical problem. While so-called 'presenting signs and symptoms' may be suggestive of one or more abnormal or pathological conditions, In most cases, conclusive evidence concerning the underlying disorder is not available.

Because he can be much more selective in the choice of therapeutic measures when attempting to treat the cause rather than its manifestations, the clinician is generally obliged to formulate a model, comprising one or more hypothesized abnormal or pathological conditions, as a basis for clinical problem solving. This process of generating and testing hypotheses with respect to unobserved - perhaps unobeervable–pathological conditions Is the essence of medical diagnosis. Because this task is one of 'problem finding' or 'problem formulation,' not 'problem solving' as this term is generally construed, It is not clear that the solution-oriented heuristic procedures of artificial intelligence have any bearing on this essential aspect of the medical task environment.

In certain apeclal cases, where the diagnostic task can be taken to be that of identifying the presence of a single clinical problem, any of a host of deterministic or probabilistic 'recognizers' can be employed to deal with the problem. (See for example: Ledley and Lusted (2), Gorry and Barnett (3), Blelch (4).) Such methods are of very limited scope, however, as they rest on the assumption that the clinical cases to be analyzed are 'pure' - i.e., there are no erroneous data (sometimes called red-herrings) to be dealt with, and only one cllnico-pathological diagnosis to be discerned in esch case.

In the real world of internal medicine, these assumptions are rarely satisfied. Nonetheless, the skilled clinician manages to sort out the facts of s case, disregarding some, while on the basis of others formulating a clinical picture that may comprise a number of interrelated or distinct disease entities.

In what follows, we describe and illustrate a model of the diagnostic process (called DIALOG, for DIAgnostic LOGic) that was designed to mimic the data structures and reasoning processes of the skilled clinician. The process employs a novel heuristic procedure for ranking and partitioning of diagnostic Inferences, which enables construction of clinically relevant diagnostic models even in the presence of erroneous and misleading data. Using an extensive medical data base which encompasses approximately fifty percent of the major diseases of Internal medicine, the DIALOG procedure has accurately analyzed many complex, real-world clinical problems, Involving aa many aa five distinct disease entities.

In the following sections, we describe first the logic of the diagnostic process and the structure of medical knowledge that underlies the DIALOG data base. There follows s description of the DIALOG procedure, and an illustration of the system's behavior.

1.2. The Logic of Diagnosis

Given a set of observations, which are taken to be consequences of some unknown cause, the task of diagnosis is to devise reasonably cogent hypotheses which would be sufficient to enable derivation, via deductive inference, of the observed consequences. This method of inference, sometimes referred to as abduction, or the method of hypothesis (5, 6), requires the design of procedures and data structures to deal with the following issues:

1. Observations must be able to 'trigger' or evoke hypotheses of disease entities with which they are associated. Since certain findings may suggest one disease more strongly than another, It is necessary to provide for some measure of strength of association between observations and the hypotheses they evoke.

2. Hypotheses must be able to generate expectations concerning likely consequences, which may be posed as questions regarding additional observations (or laboratory 'studies perhaps) in order to 'test' the hypotheses. This suggests the need for a second set of associations - those relating disease entitles to the possible consequences thereof; here, a useful measure of the degree of association is the frequency of occurrence of a given finding in a given disease.

3. Since the business of hypothesis formation is an Inexact process (some would call it an art), it is necessary to provide some means for deciding among contending hypotheses. This suggests that the candidates be ranked on some basis, and that criteria be established for deciding when the weight of evidence is sufficient to permit reasonably confident judgments to be rendered.

4. The final and perhaps most important consideration is that some means must be developed to group hypotheses Into mutually exclusive subsets corresponding to coherent problem areas. This is essential in order to be able to deal with cases where more than one disease may be present, and hence more than one hypothesis correct. The discriminating procedure sketched above must not be forced to choose between two complementary hypotheses; somehow. It must focus on those evoked hypotheses that constitute alternatives to one another, not complements.

In section 2, we show how each of these issues is addressed in the DIALOG model. Before pursuing these matters further, however, we turn our attention next to a consideration of the characteristics of the knowledge base underlying internal medicine.

1.3. Structure of Medical Knowledge

There are three significant relations that we have discerned requiring representation in the knowledge base for internal medicine. There are the two relations described above, which may be expressed as:

1) EVOKES (M, D) denoting the evocative association which manifestation M has for disease D, also

2) MANIFEST (D, M) which is the reverse association between disease D and its manifestation M.

In addition, a number of investigators (7,8) including ourselves, have independently arrived at the conclusion that a hierarchy of disease categories constitutes a significant organizing principle that should be represented in the information structure of internal medicine. Thus we define the relation:

3) FORM-OF (Dl, Dl.l) which asserts that disease (or category) Dl.l is a form of disease category Dl. For example, Dl.l might stand for hepatocellular disease and Dl for liver disease. Then further elaboration of this region of the hierarchy would eventually lead to hepatitis A, infectuous mononucleosis, and so forth.

The reason that FORM-OF appears to be a significant relation in internal medicine is the economy of representation (and processing) made possible by the following sort of quantification:

Vd(FORM-OF(Dl,d)=»MANIFEST(d,Ml))

expressing the common situation wherein all diseases of a certain class can give rise to the same manifestation.

Because of the enormity of the data base, and because of the need to qualify the EVOKES and MANIFEST relations with respect to strength of associations, it is Impractical to express the axioms of Internal medicine by means of verbal expressions, as above. Instead, we use an explicit network to encode the hierarchy of disease categories and the related sets of associations, as Illustrated in Fig. 1. Here, the solid lines denote the FORM-OF relation; dashed lines (with upward pointers) express EVOKES; and dotted lines (with downward pointers) stand for MANIFEST. The quantified concept that 'all diseases of a certain type share a common manifestation* is denoted in this network by MANIFEST and lwnvFR HnV« »t-rnrhi>d to non-terminal nodes.

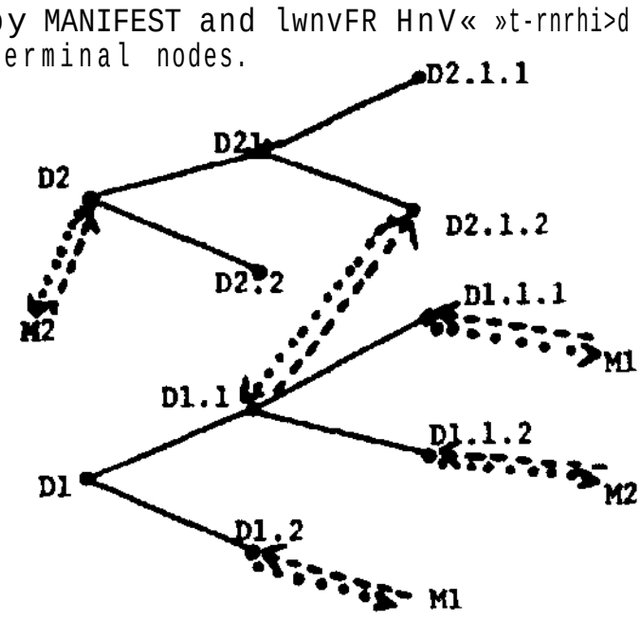

Fig. 1
Portion of Disease hierarchy with representati of EVOKES(---) and MANIFEST(...) relations.

Note that the MANIFEST and EVOKES relations can exlat between nodes of the network as well as between nodes and manifestations; this reflects the fact that one disease may be caused by, and hence be a manifestation of another.

In the following section, we give a somewhat more detailed account of the construction and interpretation of the DIALOG information structure.

2. The DIALOG System

2.1. Development of the Data Base

In constructing the disease hierarchy, successive general areas of internal medicine such as liver disease and heart disease are selected. Subcategories of each general area are chosen on the basis of similarity of pathogenetic mechanism; thus, members of a subcategory share in good part a common mode of clinical presentation. Further subdivision of subcategories is done until finally the level of individual diagnoses is reached.

Once the superstructure is completed, the appropriate manifestations are entered under each diagnostic level (terminal) node of the hierarchy. The following data are obtained: (1) a list of all manifestations of each terminal node; (2) an estimate of L(Di/Ma): the likelihood that If manifestation Ma is observed in a patient, diagnosis Dl is its cause (relative to all other causes of Ma). Estimates of this "evoking strength" are given on a 0-5 scale, with 0 indicating that the manifestation is too nonspecific to draw any diagnostic conclusions, and 5 indicating that Ms is pathognomonic for Dl; (3) an estimate of F(Mb/Dj), the frequency with which patients with proven diagnosis Dj will display Mb as a manifestation of that dlaease. A 1 to 5 scale is used, with 1 indicating Mb occurs rarely in Dj and 5 indicating Mb la a sine qua non for Dj.

In addition to listing historical findings, symptoms, signs, and laboratory data as manifestations of a given disease, diagnoses or syndromes themselves can be listed. For example, portal hypertension would be listed under Laennec's cirrhosis; secondary neoplasm of the liver under adenocarcinoma of the colon; and congestive heart failure under aortic insufficiency. These are examples of the unidirectional "causal link": D_1 is a manifestation *of* D_2 if &2 [lfl] a known cause of D^.

Once all manifestations are entered for each of the diagnosis-level nodes, a program Is invoked to carry out the 'generalization' process. For each node of the disease hierarchy, the Intersection of the manifestation lists of its subnodes is determined. The result is a list of the manifestations common to all the subnodes of a given node; it is thus the manifestation list for that higher node. By this method, jaundice becomes a manifestation of hepatocellular injury (and of cholestasis and others), and the presence of a markedly Increased alkaline phosphatase indicates that a certain category of cholestatic diseases is to be considered. Further discriminating information can be obtained only by examining the manifestation lists of subnodes In the class. This process allows the diagnostic program to construct the patient's differential diagnosis on as general a level as Is possible; thus, the program can attempt to rule out general classes of disease before working on individual diagnoses.

The final process In constructing the data baae is to enter data about each manifestation. H>at important in this area are the TYPE of the manifestation (history, symptom, sign, laboratory data, diagnosis) and the IMPORT (an Integer from one to five). The TYPE allows the diagnostic program to work on the least expensive, least dangerous differential points first (history, symptoms, signs) before going on to the more costly Items (there are 3 levels of laboratory data based on complexity, cost and danger to the patient). The IMPORT is an index of how readily an observed manifestation can be Ignored, I.e., to what degree it could be considered a "red herring." A history of shellfish ingestion has an IMPORT of 1 and ia easily Ignored; a liver biopsy showing caseating granulomas has an import of 5 and must be explained by one of the final diagnoses. Various other properties are encoded for each manifestation.

2.2. The Logic of DIALOG

The DIALOG program has been designed to mimic the problem-solving procedures of the clinician. The program begins by accepting any given sequence of presenting manifestations of illness. It then Inquires about relevant historical items, symptoms, signs, and laboratory data (proceeding via the TYPEs, exhauatlng the useful questions in each TYPE before going on *to* a more coetly TYPE).

2.2.1. Case Analysis: Entry of Presenting Manifestations (Phase I)

As each observed manifestation is entered, nodes of the disease hierarchy 'evoked by' that manifestation are processed, as follows:

(a) if the indicated node has not previously been considered - I.e., If no manifestation has previously caused this node to be evoked - a new hypothetical model is created that reflects the explanatory power of the newly considered dlsesse. Each disease model consists of four parts:

(1) a list of all manifestations that have been observed but which cannot be explained by this particular disease. This list is referred to as the 'shelf of that model.

(2) a list of all observed manifestations, along with associated evoking strengths, that are consistent with this disease.

(3) a list of all manifeatations that would ordinarily be expected to occur (with reasonable frequency) in the presence of this disease, but which have been found absent.

(4) a list of all other manifestations that are consistent with the newly evoked node, but about which nothing is yet known.

(b) If the node had previously been evoked, then all associated sublists of the model are simply updeted to reflect the new observation.

2.2.2. Case Analysis: The Interrogative Phase of DIALOG (Phase II)

Having recorded all of the Initial input data, the system proceeds to weigh.the evidence for and against each hypothesis on the evoked list. The weight assigned to each of these is determined by the following factors:

(a) counting in favor of each model is a factor proportional to the combined evoking strengths of all observed manifestations that It explains.

(b) counting against a model are two factors: data not explained are weighed In proportion to their IMPORT; data expected but found absent in the patient are weighed in proportion to their frequence of occurrence In the considered disease.

(c) a *bonus' Is awarded to those models causally linked with disease nodes that have already been confirmed.

The set of evoked hypotheses, ordered on the basis of weights computed as above, is next processed to determine which of several modes of analysis Is to be pursued. A partitioning process is utilised that employs the following concept of dominance: hypothetical model A dominates model B if the net shelf of A (i.e., the shelf minus those Items explained by previously confirmed diagnoses) is a subset of the net shelf of B. Each member M^ of the evoked models list is compared with the top-ranked model T. If Mi either dominates or is dominated by T, It is placed on the 'considered' list; otherwise it is pieced on the deferred list which is temporarily Bet aside.

This partitioning heuristic has the effect of grouping with the top ranked model those diegnosee that may reasonably be considered mutually exclusive alternatives to It. Since the combination of any of these models with the top-ranked model would add nothing to the explanatory power of the Individual models taken separately, the diagnostic process reduces (for the moment) to a discrimination among these alternatives.

Once a 'considered' list Is selected, only those diagnoses within a fixed range of the top-ranked model are used to determine which mode of questioning is to be used. When this reduced list contains five or more models, 'RULEOUT' mode Is used. 'RULEOUT* asks about manifestations with very high frequency of occurrence in the diseases being processed. Such questions stand a good chance of eliminating one or more of the considered models. The level of questions asked Is incremented via the TYPE, so that inexpensive Items are asked first. Because of the high cost associated with the acquisition of laboratory data, 'RULEOUT' mode is not used when the TYPE of questioning has reached the level of laboratory procedures. Instead, the key word 'HARROW' Is printed out, and the field of considered diagnoses is artificially narrowed so that 'DISCRIMINATE' mode can be used, which normally applies only when the reduced considered list contains from two to four models. In this mode, the top two diagnoses are selected for discrimination; Items that count heavily for one model while counting heavily against the other are the desiderata for questioning. Finally, if the reduced considered list contains only one model, 'PURSUING' mode is used. Questions are selected that are thought to have a good chance of being 'clinchers.' Manifestations which have a strong evoking strength with respect to the considered model are asked. The system continues In 'PURSUING' mode until either the Initial spread between the two top models has reached criterion, or until the spread has been reduced to the point that the top node no longer stands alone on the major liat. In the former case, the system 'CONCLUDE's that the considered disease is present; In the letter, processing reverts to the 'DISCRIMINATE' mode.

In each mode, a small number of questions are selected and asked. The responses to the set of queries are proceesed In a manner essentially the same as described for Phase I: new nodes are evoked, old nodes updated, etc. A new ranking of all evoked models Is determined; the processes of partitioning, mode selection, end Interrogation are then repeated as often as necessary. Whenever the presence of a particular disease Is concluded, the list of manifestations explained by that disease is deleted from further consideration; diseases causally related to the confirmed diagnosis are given appropriate bonus scores (dependent on evoking strength and frequency of causal relationship). Phase II processing is then repeated In an effort to discover and confirm additional problems.

2.3. DIALOG Case Analysis: An Example

Due to space conetrainte, it la not possible to include a complete protocol of the dialog that took place in analyzing the case discussed below. Rather than include all questions and answers that were exchenged between the physician and machine, we have chosen to highlight the key decision making points of the process. Anyone desiring further information may contact the authors for complete protocols of this and other cases analyzed by DIALOG.

Before turning to the substantive aspects of the transcript, some explanation of the format of the protocol is in order.

As many data aa desired can be entered initially, and the order of entry does not matter. The program prints an asterisk (*) when requesting input, and will continue to do so until the respondent types FINIS to signal that he has finished.

After processing the initial data, the program

prints out its current list of 'considered nodes,' prefaced by one of the keywords (RULEOUT, NARROW, DISCRIMINATE, PURSUING, CONCLUDE) discussed in the preceding section. If any date are not explained by the top-ranked modal, they are displayed with the notation 'DISREGARDING.'

The program then solicits additional information by one of two question types. If it displays the name of a manifestation followed by a question mark, it expects a response of YES, NO, or NA (meaning 'not available'). When it asks 'PLEASE ENTER FINDINGS oF...,' Information is being requested concerning a group of manifestations. In response, either all information about the inquired group of manifestations can be entered, or else 'NO' can be typed in order to learn by the next printout what specific question in the group the program had been considering. If the user types In a manifestation name in response to the question, the program will continue to request input (by displaying a '*') until the user types FINIS.

This case demonstrates how the program manages the problem of multiple diagnoses in the same patient. The diagnostic problem ia easier when more than one diagnosis are oausally related, but the machine oan diagnose several completely unrelated conditions. In this aaee, four of the abnormal states–pulmonary emphysema, congestive heart failure$_t$ congested liver, and the transudative pleural effusion are interrelated.

***(DOCTOR)**

SYMPTOMS?

***AGE-GTR-THAN-50**

***SEX-MALE**

***CIGARETTE-SMOKING-HX**

***HEART-FAILURE-CONGESTIVE-HX**

***WHEEZING-HX**

***CHEST-PAIN-ANTERIOR-HX**

***DYSPNEA**

***DYSPNEA-EXERTIONAL**

***ORTHOPNEA**

***COUGH**

***ANOREXIA**

***EDEMA-SUBCUTANEOUS**

***BREATHING-ACCESSORY-MUSCLES-USED**

***RALES-LOCALIZED**

***CHEST-PERCUSSION-HYPERRESONANT-BILATERAL**

***LIVER-ENLARGED**

***HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP**

***HEART-SOUNDS-S3-LEFT-VENTRICULAR-GALLOP**

***SPUTUM-MUCOID**

***BREATHING-PROLONGED-EXPIRATORY-PHASE**

***BREATHING-INTERCOSTAL-RETRACTIONS-INSPIRATORY**

***CHEST-PERCUSSION-DULL-LOCALIZED**

***BREATH-SOUNDS-DECREASED-LOCALIZED**

***HEART-MURMUR-APICAL-SYSTOLIC-EJECTION**

***PRESSURE-VENOUS-INCREASED**

***HEART-XRAY-LEFT-VENTRICULAR-ENLARGEMENT**

***HEART-XRAY-RIGHT-VENTRICULAR-ENLARGEMENT**

***TUBERCULIN-TEST-NEGATIVE**

***FINIS**

After the primary data are entered, the program already has adequate evidence to diagnose congestive heart failure.

(DISREGARDING: CHEST-PERCUSSION-HYPERRESONANT-BILATERAL HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP BREATHING-PROLONGED-EXPIRATORY-PHASE BREATHING-INTERCOSTAL-RETRACTIONS-INSPIRATORY CHEST-PERCUSSION-DULL-LOCALIZED BREATH-SOUNDS-DECREASED-LOCALIZED HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(CONCLUDE: HEART-FAILURE-CONGESTIVE)

The machine then tackles the next problem - the one about which it has the most conclusive information. That problem is the differentiation of pulmonary emphysema from bronchial asthma. This is accomplished with a modest amount of simple clinical information.

(DISREGARDING: HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP CHEST-PERCUSSION-DULL-LOCALIZED BREATH-SOUNDS-DECREASED-LOCALIZED HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(DISCRIMINATE: PULMONARY-EMPHYSEMA BRONCHIAL-ASTHMA)

PLEASE ENTER FINDINGS OF EYE-OPHTHALMOSCOPY

.
.
.

(DISREGARDING: HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP CHEST-PERCUSSION-

DULL-LOCALIZED BREATH-SOUNDS-DECREASED-LOCALIZED HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(PURSUING: PULMONARY-EMPHYSEMA)

PLEASE ENTER FINDINGS OF CHEST-PERCUSSION-PULMONARY

.
.
.

(DISREGARDING: HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP CHEST-PERCUSSION-DULL-LOCALIZED BREATH-SOUNDS-DECREASED-LOCALIZED HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(CONCLUDE: PULMONARY-EMPHYSEMA)

The next problem considered is the left ventricular disease. Simple bedside observations are not productive so the program turns its attention to the problem of pneumonia versus pleural effusion. This process of switching from one aspect of the overall diagnostic problem to another is one strength of the program and is determined by the balance of evidence favoring each problem. The deferred left ventricular problem is not permanently disregarded, however, and will be returned to in due time.

(DISREGARDING: HEART-SOUNDS-P2-INCREASED CHEST-PERCUSSION-DULL-LOCALIZED BREATH-SOUNDS-DECREASED-LOCALIZED)

(RULEOUT: AORTIC-STENOSIS ACUTE-MYOCARDIAL-INFARCTION CARDIOMYOPATHY-PRIMARY IDIOPATHIC-HYPERTROPHIC-SUBAORTIC-STENOSIS AORTIC-INSUFFICIENCY)

PLEASE ENTER FINDINGS OF HEART AUSCULTATION

.
.
.

(DISREGARDING: HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(DISCRIMINATE: PLEURAL-EFFUSION-TRANSUDATE GRAM-NEGATIVE-ENTERIC-PNEUMONIA)

PLEASE ENTER FINDINGS OF CHEST-XRAY-LUNG-FIELDS

.
.
.

PLEURAL-FLUID-SPECIFIC-GRAVITY-GTR-THAN-1:013? *YES

.
.
.

(DISCRIMINATE: PLEURAL-EFFUSION-EXUDATE PLEURAL-EFFUSION-TRANSUDATE)

PLEASE ENTER FINDINGS OF PLEURAL-FLUID-SPECIAL-EXAMINATION

*PLEURAL-FLUID-PROTEIN-LESS-THAN-3-GMS-PERCENT

(DISREGARDING: PLEURAL-FLUID-SPECIFIC-GRAVITY-GTR-THAN-1:013 CHEST-XRAY-LOWER-LUNG-FIELD-DENSITY[IES] HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(PURSUING: PLEURAL-EFFUSION-TRANSUDATE)

(CONCLUDE: PLEURAL-EFFUSION-TRANSUDATE)

As would be expected, the diagnosis of "pleural-effusion-transudate" is established after thoracentesis. But note that the evidence is conflicting--the specific gravity is too high for a transudate (probably an actual laboratory error) but on learning that the protein concentration is appropriate for a transudate, the correct conclusion is drawn.

Next, the program turns its attention to the enlarged liver and in this setting has little trouble in diagnosing hepatic congestion. The alternative diagnosis, cardiac cirrhosis, is a most worthwhile consideration but little support was found for it.

(DISREGARDING: PLEURAL-FLUID-SPECIFIC-GRAVITY-GTR-THAN-1:013 CHEST-XRAY-LOWER-LUNG-FIELD-DENSITY[IES] HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(DISCRIMINATE: HEPATIC-CONGESTION CARDIAC-CIRRHOSIS)

PLEASE ENTER FINDINGS OF LIVER-FUNCTION-TESTS

.
.
.

(DISREGARDING: PLEURAL-FLUID-SPECIFIC-GRAVITY-GTR-THAN-1:013 CHEST-XRAY-LOWER-LUNG-FIELD-DENSITY[IES] HEART-SOUNDS-P2-INCREASED HEART-SOUNDS-S4-LEFT-ATRIAL-GALLOP HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(CONCLUDE: HEPATIC-CONGESTION)

Lastly, attention is returned to the left ventricular problem and, after learning of the EKG abnormalities, myocardial infarction is diagnosed.

(DISREGARDING: PLEURAL-FLUID-SPECIFIC-GRAVITY-GTR-THAN-1:013 CHEST-XRAY-LOWER-LUNG-FIELD-DENSITY[IES] HEART-SOUNDS-P2-INCREASED HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(RULEOUT: ACUTE-MYOCARDIAL-INFARCTION AORTIC-STENOSIS MITRAL-INSUFFICIENCY CARDIOMYOPATHY-SECONDARY MYOCARDITIS IDIOPATHIC-HYPERTROPHIC SUBAORTIC-STENOSIS)

CHEST-PAIN-ANTERIOR-LASTING-GTR-THAN-20-MINUTES?

***YES**

.
.
.

(DISREGARDING: PLEURAL-FLUID-SPECIFIC-GRAVITY-GTR-THAN-1:013 CHEST-XRAY-LOWER-LUNG-FIELD-DENSITY[IES] HEART-SOUNDS-P2-INCREASED HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(RULEOUT: ACUTE-MYOCARDIAL-INFARCTION CARDIOMYO-PATHY-SECONDARY CARDIOMYOPATHY-PRIMARY MYOCARDITIS-ACUTE-INFECTIOUS IDIOPATHIC-HYPERTROPHIC-SUBAORTIC-STENOSIS)

PLEASE ENTER FINDINGS OF ELECTROCARDIOGRAM

***EKG-LEFT-AXIS-DEVIATION**

***EKG-P-WAVES-TALL**

***EKG-QRS-PROLONGATION**

***EKG-ABNORMAL-Q-WAVES**

***EKG-ST-SEGMENT-ELEVATION**

.
.
.

(DISREGARDING: EKG-LEFT-AXIS-DEVIATION EKG-P-WAVES-TALL PLEURAL-FLUID-SPECIFIC-GRAVITY-GTR-THAN-1:013 CHEST-XRAY-LOWER-LUNG-FIELD-DENSITY[IES] HEART-SOUNDS-P2-INCREASED HEART-MURMUR-APICAL-SYSTOLIC-EJECTION)

(CONCLUDE: ACUTE-MYOCARDIAL-INFARCTION)

The patient from whom the data of this case are derived had died some weeks earlier and all five conclusions were verified by autopsy.

Summary

We have presented in this paper an approach to the problem of hypothesis formation, which is central to the task environment of internal medicine and to many other areas of intellectual endeavor as well. There are, of course, many other ways to attack the problem, some of which have been reviewed in Pople (9).

In our view, the most significant aspect of the DIALOG approach is the focusing heuristic that partitions hypotheses Into coherent problem areas, thereby inducing system behavior that resembles the *problem oriented' approach of the skilled clinician.

This approach would appear to have validity in any problem formulation task characterized by the potentlel concurrence of multiple entities, requiring ad hoc assembly of elementary hypotheses. We commend It to those investigating the 'Theory of Frames'(10).

Bibliography

1. Nilsson, J. J., Problem Solving Ms thods in Artificial Intelligence. New York: McGraw-Hill, 1971.

2. Ledley, R.S., and Lusted, L.B., Reasoning Foundation of Hadical Diagnosis: Symbolic Logic, Probability and Valua Theory and our Understanding of How Physicians Reason, Science. 130:9, 1959.

3. Gorry, G.A. , end Bernett, C.O., Experience with e Model of Sequential Diagnosis. Computer Biomedical Research, 1:490, Nay, 1968.

4. Bleich, H.S., Computer Evaluation of Acid-Base Disorders. J. Clin. Invest., 48: 1689, 1969.

5. Pelrce, C.S., Collected Papers of Charles Sanders Pelrce. C. Harthshorne, P. Weiss, and A. Burks (Eds.) 8 vols., Cambridge* Mass.: 1931-1958, asp. Vol. I I , pp.272-607.

6. Pople, H.E., Jr., On the Mechanisation of Abductive Logic, Proc, 3, IJCAI.

7. Wortmen, P.M., "Medical Diagnosis: An Information Processing Approach," Comput. Blomed. Res. 5, 315-328, 1972.

8. Rubin, A.D., "Hypothesis Formation and Evaluation In Medical Diagnosis," Masters Thesis, Dept. of Electrical Engineering, M.I.T., 1975.

9. Pople, H.E., Jr., Artificial-Intelligence Approaches to Computer-Based Medical Consultation, IEEE Intercon Conference, 1975.

10. Mlnsky, M. "A Framework for Representing Knowledge," in The Psychology of Computer Vision, P.H. Winston, ed., McGrav-Hlll: Mew York, 1975.

Elementary hypotheses can sometimes themselves be considered symptoms, so that EVIDENCE and EXPECTATION pointers May connect them. In Diagram 1, UTI and PYELONEPHRITIS, both elementary hypotheses, are so related; the symptoms of UTI are a eubset of those of PYELONEPHRITIS. There are also clearcut CAUSE relations betueen elementary hypotheses where the symptoms of the two diseases concerned *are* not in a subset/superset relation. In this case, the CAUSE relation is stated explicitly, as in STREP-INFECTION CAUSES AGN. Similar to CAUSE links are COMPLICATION links, as in PYELONEPHRITIS is a COMPLICATION of STONE. In addition, elementary hypotheses may or may not be ULTIMATE-ETIOLOGIES. An elementary hypothesis which is an ULTIMATE-ETIOLOGY is one which could stand alone as a diagnosis, for which a more basic cause does not have to be sought or is not known.

The system knows about several different types of findings: LAB-DATA, PHYSICAL-EXAM. SYMPTOM, FACT and FAMILY-HISTORY. The process of deciding whether or not a particular patient finding is relevant to the symptom description in the knowledge network and thus relevant to the disease hypothesis is called fittina; (see <Winograd> and <Minsky> *for* earlier uses of this term in frame theory) It requires trying to fit a particular finding-description into a sometimes more general specification. The result of this fitting process is that the finding-specification is either confirmed or disconf irmed. if the finding contains enough detail; if more information is needed to see if the finding fits the specification, the doctor often asks more questions.

The time of occurrence of the symptoms and the time relationships indicated in the knowledge net are also taken into account in fitting, as described in <Rubin>.

During the course of a diagnostic session, nodes of the data network change state with the addition of new information about the patient. while finding-specifications may be either confirmed or di sconf irmed. elementary hypotheses, because they are not directly confirmable, have a more complicated set of alternative states. when a diagnostic session starts, all elementary hypotheses are inactive: that is, no particular disease has been suggested by the patient's symptoms. As more data is presented, certain hypotheses become active by virtue of their correlation with and ability to account for the findings present. Once a hypothesis it active, it is evaluated after the addition of every finding to see how well it fits the data so far and some score is produced which represents the likelihood of that disease's being present. On the basis of this process, a hypothesis may be accepted or re iectedi in most cases, however, no definite decision will be made, but its score will be modified to reflect the effect of the new data. An accepted elementary hypothesis is one for which the evidence is sufficiently specific to rule out any other cause for the symptoms present. For example, the presence of REO-BLOOD-CELL-CASTS confirms the diagnosis of GLOMERULI TIS, making it an accepted hypothesis, but the very same finding makes SICKLE-CELL-TRAIT a rejected hypothesis Finally, consideration of an elementary hypothesis may be deferred until more supporting symptoms are known. This process helps reduce the number of concurrently-active hypotheses.

How does the magic transformation from a bunch of symptoms to a final diagnosis take place? The process seems to be divided into four steps: di sbosina. triggering, local evaluation and glojtfl assembl ino. This series of four steps is performed after the addition of each finding. A brief description of the four stages follows; the third and fourth will be discussed in more detail below.

1.Sometimes the cause of a finding i« clear when the finding is encountered; this is most often the case whan the explanation is a FACT. In such a circumstance, the doctor di sboees of the finding as a result of some already-accepted ettoJogy rather than trying to find a new explanation. For example, suppose a patient is brought into the emergency room of a city hospital after an automobile accident; if his urine contains blood, the doctor should surely attribute it to abdominal trauma, rather than considering GLOMERULITIS.

2.Triaaer ino is a process by which *an* elementary hypothesis makes the transition from the inactive to the active state. A subset of the symptoms which *are* relevant to a disease are marked as triggers. Uhen a symptom is asserted to be present in the current case, it act ivates all those elementary hypotheses for which it has been designated a trigger. For example, DYSURIA (painful urination) triggers URINARY-TRACT-INFECTION; NAUSEA by itself triggers nothing, as it is a common finding in many disorders.

3.Each elementary hypothesis has an associated local evaluation function which produces a value representative of how likely the disease is to be present given the data. Each of the hypotheses currently active is evaluated, taking the new finding into account. The evaluation done at this stage Is local in that the functions do not ask questions about the status of other

elementary hypotheses, or consider symptoms other than thoee relevant to the diseaee hypothesis being evaluated.

4.The purpose of global assemblu. the fourth stage, is to arrange the various local elementary hypotheses into a larger structure which is both coherent and adequate. The rules of coherence have to do with the wage to connect various elementary hypotheses through links like CAUSE. COMPLICATION and EVIDENCE. An adequate hypothesis is one which accounts for all the abnormal findings in a case and is the end goal of a diagnostic process.

Local Evaluation

Any hypothesis-based theory needs a method for evaluating an elementary hypothesis - in this case, a single disease or syndrome - in isolation. Correlations between symptoms and diseases are the major determinants of such evaluation. One such correlation is the conditional probability of a symptom given a disease. I have called such numbers EXPECTATIONS. In Bayesian terms they are P(5/D) (read "the probabiIity of S given 0"), where S is the symptom and 0 the particular disaase in question. These correlations, however, are all disease-centered: that is, they spring directly from the description of a disease. (lore useful diagnostic information is sumotom-centered. since a diagnosis proceeds from symptoms to diseases.

This other more sophisticatea type of information is what I have termed EVIDENCE; in Bayesian terms, it ie the conditional probability of a disease given a symptom, or P(0/S). The complexity of the transformation of information from EXPECTATIONS to EVIOENCE (see <Feller>) makes plausible the idea that part of a doctor's expertise lies in the translation of knowledge from the disease-centered mode to the symptom-centered mode.

A local scoring algorithm must take into consideration both positive and negative contributions to the current hypothesis. In general the presence of relevant symptoms (EVIDENCE) will add to the validity score of an elementary hypothesis, uhi Ie their absence (VIOLATED EXPECTATIONS) ui I I subtract from it. The presence of FEVER will add to the validity of STREP-INFECTION, while its absence will subtract. The scoring algorithm we developed takes into account both positive and negative evidence for a hypothesis and scores are normalized by being divided by their highest possible total score.

The scoring algorithm also takes into account differing strengths of properties, the inclusion of ISA links and symptoms associated uith a hierarchy of diseases, and age and sex of the patient, and operates with only four levels of EVIDENCE (SUFFICIENT.STRONC. MODERATE,UEAK) and four of EXPECTATION (NECESSARY, STRONG, MODERATE,UEAK). Oetails of the algorithm can be found in <Rubin>, but more important to the present discussion are the aspects of local evaluation which serve to reduce the number of hypotheses actively being considered at any given time.

A hypothesis may be activated by one of its triaoers. A hypothesis which is not active is not being currently considered or evaluated. This selective activation of hypotheses is one way to control the number of diseases being actively considered at any time. Notice that thie use of triggers is certainty a heuristic device, since the diagnosis for the particular case on hand may not be one of those triggered.

A second way of minimizing hypotheses is to reject unlikely ones. One method of rejection is to include in a disease's siice rejecting aumotoms whose presence precludes that disease's existence. For example, the presence of RED-BLOOO-CELL-CASTS rules out the diagnosis of SICKLE-CELL-TRAIT.

In addition, we can assign a priori probabilities to diseases. The age and sex of a patient affect this probability profoundly. Combining age, sex and disease leads to a useful number representing the probability of the disease occurring in a patient of particular age and sex. If this number is especially low, we may consider it 0 for heuristic purposes and put the hypothesis on the DEFERRED-LIST.

The theory presented so far has been a I inear one. Such a theory assumes that subparts of a problem can be treated independently and the solutions to those subproblems combined without alteration. Of course, there are interactions among symptoms which contradict the linearity hypothesis. The three examples of non-linearity described below illustrate some ways of dealing with interactions.

Both HEMATURIA and PROTEINURIA are EVIDENCE f o r GLOMERULITIS and G-U-TRACT-BLEEDING. Houever, their relative severities differ in these two hypotheses. In G-U-TRACT-BLEEOING, we expect the ratio of HEMATURIA to PROTEINURIA to be near that in uhole blood; for HEMATUIRA GROSS we expect PROTEINURIA LIGHT. In GLOMERULITIS, on the other hand, there should be relatively more PROTEINURIA than in G-U-TRACT-BLEEDING ; f o r PROTEINURIA M00ERATE, we would expect HEMATURIA MICROSCOPIC or LIGHT. The approach I have taken to this interaction is to specify for each disease or state which combinations would rule j j . QUI. Thus (AND (HEMATURIA GROSS) (PROTEINURIA LIGHT)) precludes GLOMERULITIS, while (AND (HEMATURIA LIGHT) (PROTEINURIA HEAVY)) precludes G-U-TRACT-BLEEDING.

The ASLO (anti-streptolysin-0) titer often rises several weeks after a person has had a STREP-INFECT I ON, indicating that the body is fighting the infection with antibodies. Taking PENICILLIN to combat the infection, however, often squelches the antibody response. If a doctor were actively considering STREP-INFECTION. ASLO-TITER

(RESULT NORMAL) would represent a violated expectation. An excuse is sometimes available for the absence of an expected finding; in this case PENICILLIN (STATUS TAKEN) would excuse a normal ASLO-TITER. The STREP-INFECTION hypothesis is evaluated as if ASLO-TITER uere not a relevant symptom uhen penicillin has been taken.

Two or more diseases may resemble each other in many of their crucial aspects; it is particularly important to be able to tell them apart. Besides being a possible pitfall in causing misdiagnoses, findings which are shared among diseases can also be used heuristically to avoid activating an undue number of hypotheses. Suppose diseases A *and* B share findings X, Y and Z, but are differentiated by Q's occurrence in A but not B. If X and Y are present, we can consider B but not A, provided there is also a piece of heuristic information (called a DIFFERENTIAL-DIAGNOSIS) which activates A and rejects B if Q is discovered.

Global Assembly

Host of the diagnoses at which doctors finally arrive are not represented by a single elementary hypothesis. Patients often have more than one related or even totally unrelated diseases. A final diagnosis may be GLOMERULITIS COriPLICATED-BY NEPHR0TIC-SYN0ROT1E or FGN and HYPERTENSION ESSENTIAL. Clearly we need some way to discover and specify these more complex hypotheses as well as to combine pathological states which are themselves elementary hypotheses into a larger hypothesis which postulates a single cause for alt of them. These concerns *are* handled by the aIobaI assemblu stage of processing, which puts together coherent and adequate hypotheses. The processes in this stage are described in terms of matching a pattern (a template) and performing some action on the basis of that match.

A coherent hypothesis consists of two or more elementary hypotheses joined by "coherence links" which include ISA, CAUSE, COflPLICATION-OF. 0EVEL0PS-1NT0 and EVIDENCE links. Coherent hypotheses are constructed out of already-active hypotheses and, perhaps, some inactive ones as well, which are activated in the course of constructing the larger hypothesis. Each type of coherent hypothesis can be represented as a "template" uhich is placed on the patient's data structure; present findings, active hypotheses and necessary links are the structures which must match. The action taken when a template matches consists of joining the matched components together, along with newly-activated hypotheses. Two coherent hypothesis types are described below.

Symptoms are often not connected directIy to their diseases, but to intermediately-general pathological states. It is thus important to be able to reconnect the symptoms to the actual disease; this is done via EVIDENCE-chained hypotheses. They are formed uhen two or more active elementary hypotheses have EVIDENCE chains which intersect at a single etiology. For example, if SO0IUM-RETENTI0N and ACUTE-RENAL-FAILURE were both active, we would want to unify them into a hypothesis which postulated AGN. See Diagram 2 for an illustration of a template for this type of hypothesis.

Forming a coherent hypothesis with CAUSE. COtlPLICATION-OF or DEVELOPS-INTO links may not involve activating any new elementary hypotheses at all. If two active hypotheses are connected by a CAUSE, COtlPLICATION-OF or DEVEL0PS-INT0 link, they may be joined into one composite hypothesis. More interesting is the case uhere a new elementary hypothesis must ba activated. A template for this situation is contained in Diagram 3. CLOTTING-DISORDER is activated in order to provide the link between HEMATURIA and PREGNANCY. In words, we can activate an intermediate hypothesis which has one (non-trigger) finding present and is also connected by a CAUSE, C0HPLICATI0N-0F or DEVELOPS-INTO link to an active hypothesis. The neuly-activated hypothesis provides a link between the two nodes.

Ultimately, the global assembly stage must come up with an adequate hypothesis. The primary characteristic of an adequate hypothesis is that is accounts for all the abnormalities noted, uhile maintaining as much simplicity as possible. An adequate hypothesis consists of several independent parts, each of which is a coherent hypothesis. Each component must also be an ULTIMATE-ETIOLOGY or, in the case of more complex coherent hypotheses, it must contain some ULTIMATE-ETIOLOGY. In addition, all accepted elementary hypotheses must be subsumed in the final diagnosis, either by themselves, or as part of a larger coherent hypothesis. For example, the following is an adequate hypothesis:

LGN
(DURATION (YEARS 10))
HYPERTENSION ESSENTIAL
(DURATION (YEARS 5))
FAMILY-HISTORY NEPHRITIS

Notice that the second component of this hypothesis is HYPERTENSION ESSENTIAL, not HYPERTENSION CHRONIC; this is because only HYPERTENSION ESSENTIAL is marked as an ULTIMATE-ETIOLOGY. HYPERTENSION CHRONIC is a symptom, not an explanation, while HYPERTENSION ESSENTIAL is an explanation (actually the admission that no other explanation has been found') Essentially, a process which builds adequate hypotheses must partition the symptoms into possibly non-disjoint subsets and account for each subset with some coherent hypothesis. Strategies for this partitioning (sea <Rubin>> contribute further to the process of minimizing the number of active hypotheses.

Summaru

The area of medical diagnosis has been investigated as an A.I. problem and a

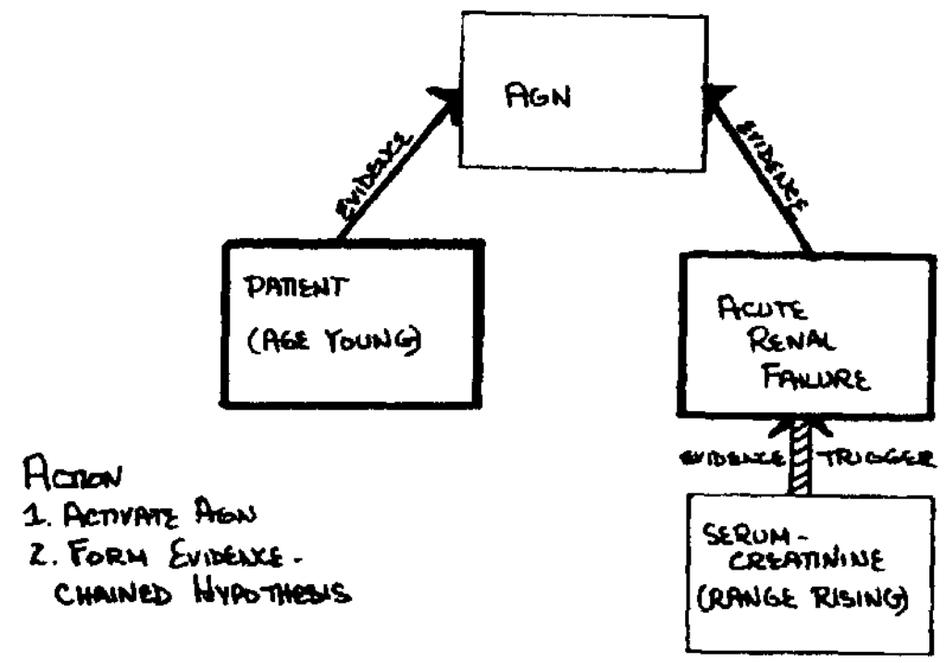

ACTION
1. ACTIVATE AGN
2. FORM EVIDENCE-CHAINED HYPOTHESIS

Diagram 2: EVIDENCE-CHAINED TEMPLATE

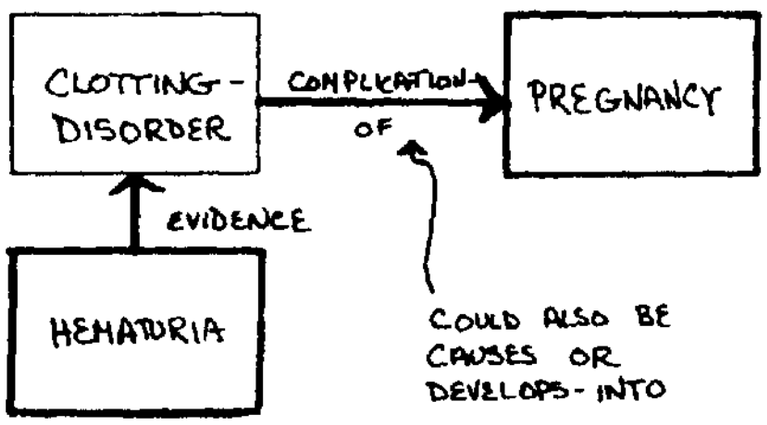

ACTION:
1. ACTIVATE CLOTTING-DISORDER
2. FORM COMPLICATION-OF
CONNECTED HYPOTHESIS

Diagram 3: GENERAL CAUSES, COMPLICATION-OF, DEVELOPS-INTO TEMPLATE

structure for medical knowledge proposed. The process postulated to act on that structure has as one of its goals the Minimization of the number of hypotheses actively considered *at* any one time. Local vs. global evaluation and linearity vs. interaction have also been studied within the context of the theory.

RefproncoB

1. Fahlman, Scott, "A Hypothesis-Frame System for Recognition Problems," WP-57, Artificial Intelligence Lab, M.I.T., Cambr idge, December, 1973.

2. Feller, Wilfiam, An Introduction to Probability Thopru and Its Applications. Vol. 1, John Uiley and Sons, Inc.: New York, 1968.

3. Gorry, G. A., ft Sufllfifi for CPWUtftT-Aided Diagnosis. TR-44, Project MAC, M. I. T., Cambridge, September 1967.

4. Gorry, G. A. et. al.. Computer Laboratory for Clinical Decision-flaking, research proposal to NIH. flarch 1974.

5. "A Framework for Representing Knowledge." in The. Psucholoau of Computer Viflion. P.H. winston, ed., McGraw-Hill: New York 197S.

6* Rubin, Andee, "Hypothesis Formation and Evaluation in Medical Diagnosis,"

AI-TR-316, Artificial Intelligence Lab, n.I.T., Cambridge, January 1975.

7, winograd, Terry, "Frame Representations and the Procedural/ Declarative Controversy," in Bobrow and Collins (eds.), forthcoming.