Mini-Max Procedure versus Bayes Formula to Calculate Probability of a Medical Diagnosis

CARLOS FEDER, MD, Internal Medicine; TOMAS FEDER, PhD, Computer Sciences

In the 18th century, Thomas Bayes, a theologian and mathematician, proposed his formula for calculating conditional probabilities, which was posthumously published in 1763. To our knowledge, Ledley and Lusted [1], in 1959, committed the original sin to apply this formula to calculate the probability of a specific diagnosis, given the manifestations—clinical data—of a disease. After that, Bayes formula gained widespread dissemination in the field of medical diagnosis. This problem became compounded by the development of Bayesian and neural networks. The inaccuracy of Bayes formula for calculating the posterior probability of a diagnosis given clinical data that are not independent has been recognized by many researchers. Their excuse for using Bayes despite this realization is that they consider the introduced error not to substantially affect the final result [2]. Our novel mini-max procedure, meant for computerized diagnosis, calculates the probability of a specific diagnosis, processing clinical data present—favoring the diagnosis—and clinical data absent—disfavoring the diagnosis. It is more accurate than Bayes formula and has remarkable additional properties, which enable diagnosis of concurrent diseases and recommendation of the best cost-benefit *set* of clinical data further to investigate to reach end of diagnostic quest through the most efficient pathway.

Bayes formula is accurate only when three conditions are fulfilled [3] [4] [5]:

1. Clinical data used for calculation of the conditional probability of a diagnosis must be **independent**: that is, a specific clinical datum—symptom, physical sign, test, or procedure result—should neither favor nor disfavor any other clinical datum of the same disease. In other words, the probability that one clinical datum is manifested by a specific disease, should not depend on the presence of another clinical datum. This is not true in actual clinical cases, where clinical data result from a chain of reactions that originate in a common cause or lesion and are necessarily related. These clinical data configure syndromes that by definition are associations of related clinical data (*e.g.*, jaundice, increased blood bilirubin, and dark urine.)

2. The diseases must be **incompatible,** which means that clinical data justified by one disease cannot be justified by another disease. When concurrent diseases occur, some clinical data may be caused by more than one of them. Because Bayes formula is only capable to calculate probabilities of competing diagnosis, which are incompatible because only one can become a final diagnosis, it is unsuitable to handle concurrent diseases. To solve the problems of independence and incompatibility, so-called Bayesian networks have been devised, but their application to diagnostic algorithms is excessively complicated and hard to compute. We created the mini-max procedure (to be explained later), which identifies concurrent diagnoses despite manifesting shared clinical data, and calculates the probability (P) of such diagnoses independently, circumventing the *incompatibility* condition of Bayes' formula.

3. The number of diseases processed by Bayes must be **exhaustive**: all known diseases must be included in the knowledge base, enabling inclusion in the denominator of the formula of all corresponding diagnoses that manifest a clinical datum. If this condition is violated, some clinical datum originated by a disease not included in the formula will distort the calculated result. Accordingly, the calculated probability of the diagnosis under consideration will be incorrect and will adversely affect the differential diagnosis. This is why computer programs based on Bayes formula but circumscribed to

a restricted area of diseases—such as congenital cardiopathies or nephropathies—are inherently inaccurate.

Our novel mini-max procedure is able to calculate the probability of a specific diagnosis, processing simultaneously clinical data present—favoring the diagnosis—and clinical data absent—disfavoring the diagnosis. It is more accurate than Bayes formula and has remarkable additional properties:

- It identifies and processes concurrent diseases.
- It precludes reduction of great probability of a confirmed diagnosis (strongly supported by clinical data present) by some relatively unimportant clinical datum absent.
- It facilitates the recommendation, at each diagnostic stage, of the best cost-benefit clinical datum next to investigate in the patient, and even a **set** of such clinical data to be investigated simultaneously.

The mini-max procedure became the core of our diagnostic algorithm, which is discussed in great detail in the book of one of us, *Computerized Medical Diagnosis: A Novel Solution to an Old Problem* [6].

To understand how our mini-max procedure works, we must first review the principles upon which it is based.

PRINCIPLES OF OUR DIAGNOSTIC COMPUTER PROGRAM

DISEASE MODEL

A **disease model** (Table 1), as defined in this study, is an abstract concept that comprises all clinical data that can be manifested by all patients with a specific disease. A single patient typically never manifests all clinical data that the disease potentially can provoke. Integration of a specific disease model with all of its possible manifestations requires statistical study of a large patient population. *Each clinical form, stage, and degree of disease has its own disease model*. Because health, death, and iatrogenic diseases are diagnoses that must be established clinically, the corresponding disease models must also be created.

Each disease model is stored in the knowledge base, listing all the clinical data that a given disease potentially can manifest.

DISEASE MODEL FOR ACUTE APPENDICITIS				
Symptoms	**S**	**PP value**	**Cost**	**Masked by**
Anorexia	0.93	0.20	none	
Pain in right lower abdomen	0.75	0.17	none	‡ (analgesics, antibiotics)
Vomiting	0.66	0.16	none	‡ (antiemetics)
Nausea	0.64	0.15	none	
Fever/chills	0.29	0.07	none	‡ (antipyretics, antibiotics)
Constipation	0.11	0.02	none	
⋮	⋮	⋮	⋮	
Signs				
Rebound tenderness	0.86	0.18	none	‡ (analgesics, antibiotics)
Fever (>37.5 C)	0.36	0.08	none	‡ (antitermics, antibiotics)
⋮	⋮	⋮	⋮	
Laboratory				
Increased white blood cell count	0.96	0.21	small	
Albumin in urine	0.19	0.04	small	
⋮	⋮	⋮	⋮	
Abdominal ultrasound				
Swollen appendix or abscess	0.60	0.14	intermediate	
⋮	⋮	⋮	⋮	
Laparotomy finding	1.00	1.00	great	

TABLE 1. Example of disease model; S, sensitivity; PP value, positive predictive value; cost of obtaining the clinical datum; ‡, interaction identifier. The numeric values for S and PP value in the above examples were not obtained from actual statistics or calculations.

DISREGARD QUALITIES OF CLINICAL DATA

Clinical data, especially subjective symptoms, typically have diverse non-exclusive qualities. For example, chest pain of angina pectoris typically is retrosternal, radiating to the neck, jaw, and upper extremities; is oppressive, lasting only a few minutes; is exertion related and relieved by nitroglycerine. Some authors confer values to these pain qualities, their chronology, and their evolution. This is correct, when such qualities, powerfully suggest a diagnosis. Nevertheless, our algorithm purposely *does not consider such clinical data qualities*; we believe that computation of clinical datum qualities is not critical for calculation of probability of a diagnosis. Reasons are: clinical data qualities and chronology are subjective and widely variable; chest pain of angina pectoris sometimes is mild, referred to the upper abdomen, not radiating, is burning, or even absent in patients with diabetes. Accordingly, these qualities may not be reliable. Anxious or hypochondriac patients can imagine such qualities. To confirm angina pectoris, more reliable tests, such as stress ECG and sometimes angiogram are needed, which provide clinical data with greater supporting value that anyway will supersede the oppressive quality of chest pain that has a lesser supporting value. Disregarding these unreliable qualities simplifies the diagnostic process without losing accuracy. It would be difficult, if not impossible, to determine the sensitivity, necessary to calculate the supporting positive predictive value of each diverse quality that thousand of known clinical data and diseases can manifest. In the especial case where the quality of a clinical datum is essential, such as the case of prolonged retrosternal pain for myocardial infarction, this quality can be included as a separate clinical datum in the corresponding disease model.

DISREGARD DISEASE PREVALENCE

Prevalence of a disease is the number of existing cases in a given population at a specific time. Prevalence statistics are of epidemiological importance. However, it may be harmful to include prevalence values when calculating the probability of a patient having a rare disease. This happens because the small prevalence value for such a rare disease could considerably reduce the probability of the corresponding diagnosis, causing it to be improperly ruled out. If a patient has a disease afflicting only one in a million persons, the probability of that diagnosis would be very small, but for him it represents one hundred percent. A perfect program should diagnose every possible disease, including those that are rare. After all, we do not need a computer to diagnose a common cold during an epidemic. Furthermore, accurate epidemiological information is difficult to obtain because many disease cases remain unreported. Accordingly, our diagnostic algorithm purposely does *not* take prevalence into account; this is equivalent to assuming that all diseases occur with the same probability.

INDICES AND IDENTIFIERS OF CLINICAL DATA

In our knowledge base, three **indices** are associated with each clinical datum: *sensitivity*, *positive predictive value*, and *cost*. Two **identifiers** are associated with selected clinical data: *risk* and *interaction identifiers*.

Indices

1. *Sensitivity* (S)

Sensitivity is defined as the conditional probability P of a clinical datum C, given a disease D:

$$S = P(C|D) \tag{1}$$

Where: S = sensitivity of clinical datum C for disease D

A practical way to calculate S of a specific clinical datum for a given disease is to determine statistically the fraction of patients afflicted by this disease who manifest the clinical datum:

$$\text{Sensitivity (S)} = \frac{\text{Number of disease cases manifesting the clinical datum}}{\text{Total number of disease cases}} \tag{2}$$

Sensitivity can be expressed either as a decimal (*e.g.*, 0.30), or as a percentage (*e.g.*, 30%.)

A given clinical datum can be manifested by more than one disease. Accordingly, both the clinical datum and the disease that manifests it, determine the value of S. This value is stored in the knowledge base linked to the corresponding clinical datum and disease model.

If the numerator and denominator of equation 2 are equal, the sensitivity (S) of the datum will equal 1, which is unlikely, as it requires that all clinical cases so far reviewed manifested the clinical datum. Otherwise, the numerator will always be smaller than the denominator, S will be less than 1, and an additional clinical case will increase S if the clinical datum is present, or reduce it if the datum is absent. Accordingly, the computer recalculates the sensitivity of each clinical datum each time the knowledge base is updated with new cases. The greater the number of cases analyzed, the greater the accuracy of S.

If a clinical datum never is manifested by a specific disease, its S equals 0 for this disease. When sufficient number of cases have been reviewed, the disease model will include all the clinical data this disease can potentially manifest, and the sensitivities will approach their true values.

2. *Positive predictive value* (PP value)

Next, it is necessary to define an index that represents the strength with which each clinical datum present in the patient supports a specific diagnosis. S cannot be used directly because it only expresses how frequently a disease manifests a clinical datum, but not how often it occurs with other diseases. We consider that positive predictive value (PP value) best accomplishes this function. PP value is defined as the conditional probability P of a disease D, given a clinical datum present C:

$$\text{PP value} = P(D|C) \tag{3}$$

The greater the PP value of a clinical datum, the stronger this clinical datum favors the corresponding diagnosis.

PP value of a clinical datum for a specific diagnosis can be calculated based on the sensitivitiy of this clinical datum for each of all the diagnoses that manifest it.

Let's start with Bayes formula, which calculates conditional probabilities:

$$P(D_i|C) = \frac{P(D_i)\,P(C|D_i)}{P(D_1)\,P(C|D_1) + \ldots + P(D_i)\,P(C|D_i) + \ldots + P(D_n)\,P(C|D_n)} \tag{4}$$

Where: $P(D_i)$ = probability of disease D_i; also called *prior probability* because it is the probability of the disease *before* considering clinical datum C

$P(D_i|C)$ = probability of disease D_i, given specific clinical datum C; also called *posterior probability* of the disease because it is the result of the equation *after* considering clinical datum C

$D_1 \ldots D_n$ = all diseases that manifest clinical datum C, including D_i

$P(C|D)$ = probability of clinical datum C, given a disease D; it equals the sensitivity (S) of the clinical datum for this disease D: $P(C|D) = S$. This is valid for any disease ($D_1 \ldots D_n$) that manifests clinical datum C

We explained earlier the reasons why we purposely do not take into account prevalence of diseases, called prior probability of diseases here in Bayesian context. This is equivalent to assuming that all diseases have the same prior probability [P (D)]; accordingly, we can simplify equation 4 by deleting the prior probability of all diseases [$P(D_1) \ldots P(D_i) \ldots P(D_n)$]. Then, if we replace P (D|C) with PP value (according to equation 3), and P (C|D) with S (according to equation 1), we obtain the following equation:

$$\text{PP value}_i = \frac{S_i}{S_1 + \ldots + S_i + \ldots + S_n} \tag{5}$$

Where: PP value_i = positive predictive value of the clinical datum for the disease i under consideration

S_i = sensitivity of the clinical datum for the disease i under consideration

$S_1 \ldots S_n$ = sensitivities of the *same clinical datum* for corresponding diseases*

* "Corresponding diseases" could either refer only to diseases that manifest the clinical datum, or alternatively to all known diseases. For either of these alternatives, the result will be identical, because S of a clinical datum for a disease that never manifests such datum is zero. Adding zeros to the value of the denominator established by S of the diseases that manifest the clinical datum will neither change the value of the denominator nor the result of the equation.

Equation 5 shows that S_i (numerator of the right member) and PP $value_i$ are directly proportional.

This equation is convenient because it expresses PP value as a function of sensitivities (S), being S the cornerstone of our algorithm.

PP value = 1 when the clinical datum is manifested only by the disease under consideration (that is, when for all other diseases S = 0); conversely, PP value approaches 0 when the clinical datum is always manifested in all other diseases (that is, when for all other diseases S = 1), a theoretical situation. In remaining situations, PP value takes an intermediate value between 0 and 1.

PP value quantifies how characteristic or exclusive a clinical datum is for a specific disease or diagnosis. According to equation 5, the fewer the number of diseases that manifest a given clinical datum (number of S terms in the denominator) and the fewer this clinical datum is manifested by each of these diseases (the smaller each S value in the denominator), the greater the PP value of the clinical datum and the probability of the specific diagnosis or disease. For example, the presence of *Mycobacterium tuberculosis* in sputum is pathognomonic of pulmonary tuberculosis because no other disease manifests this clinical datum; accordingly PP value = 1. We believe that *PP value is the most accurate index of how strongly a clinical datum supports a diagnosis or disease*.

Calculated PP values are linked to the corresponding clinical data in the disease models. Because PP values are based on statistically established sensitivities stored in the knowledge base, they do not depend on specific clinical cases, and therefore can be pre-calculated (before the diagnostic program is applied to actual clinical cases) saving computing real time. These values remain fixed unless new disease models are added to the knowledge base or revised statistics change the values of the sensitivities upon which PP values are based. Should such changes occur as a result of an occasional update, all PP values must be recalculated.

Definition of PP value based on equation 5 is rational, simple, accurate, practical, and novel.

For the diagnostic algorithm to function properly, the knowledge base must include **all** currently known diseases. Should a disease be omitted, its disease model will not be created and the sensitivities of the corresponding clinical data will not be available and will miss from the denominator in equation 5. Consequently, the calculation of the PP value of a clinical datum for any disease will be inaccurate; still worse, the excluded disease will never be included in the differential diagnosis.

Equation 5 that calculates PP values of clinical data, derives from Bayes formula, when deleting the prior probabilities. This equation does not violate the aforementioned conditions because it includes in its denominator S of all known diseases capable of manifesting the clinical datum, complying with the *exhaustive* condition. These S refer to clinical data pertaining to unrelated disease models (as opposed to a specific patient) complying with the *independent* condition. Because of the *incompatible* condition of Bayes formula, our algorithm processes concurrent diseases in a way that does not violate this condition, as explained later.

3. *Cost*

Cost is another index that is attached to each clinical datum. In our context it involves not only expense, but also risk and discomfort resulting from the required test or procedure. *Expense* is quantifiable in dollars or any other currency. *Risk* can be statistically quantified by outcomes of the procedure, although it also depends on operator skill. *Discomfort* is a subjective feeling that depends in part on the invasiveness of the procedure and in part on patient apprehension, although the latter can be controlled with sedation or anesthesia. Discomfort cannot be expressed as an exact numerical value, but only can be assigned an estimated qualitative level such as none, small, intermediate, or great. Expense, risk, and discomfort—like apples and oranges—cannot be arithmetically combined into an exact overall cost. However, expense and risk can be qualitatively expressed in levels similar to discomfort, to make the latter comparable to the former two. It is practical to consider the maximum qualitative level of expense, risk, and discomfort, as representative of overall cost level.

$$\text{Cost} = \max(\text{expense, risk, discomfort})$$

Because cost does not participate in the calculation of the probability of diagnoses, its inexactness is not critical; it is considered only when selecting the most suitable clinical datum next to investigate in the patient, as explained later.

We assign to each clinical datum one of four overall cost categories: no cost (clinical data typically obtained through medical history and physical examination), small cost (*e.g.*, obtained through routine laboratory analysis, ECG, and other ancillary studies), intermediate cost (*e.g.*, colonoscopy, lymph node excision biopsy), and great cost (*e.g.*, liver biopsy, laparotomy.)

Cost must be compared to the *benefit* expected to result from acquiring a clinical datum. Benefit has two components: a quantitative component and a qualitative component. The *quantitative component* depends on the positive predictive value (PP value) and sensitivity (S) of the clinical datum, which in turn determine the probability (P) of the corresponding diagnoses. PP value of a clinical datum present in a patient tends to increase the P of the corresponding diagnoses; S of a clinical datum absent tends to reduce the P of the corresponding diagnoses. The clinical datum that has the greatest PP value or the greatest S will result in the greatest benefit because it increments the difference between the probability of the most likely diagnoses and the probabilities of those less likely, which are ultimately eliminated from the differential diagnosis list. The magnitude of the increment of the aforementioned difference of diagnostic probabilities quantifies the benefit of the clinical datum that produces it. The quantitative component of benefit can be determined before actually investigating a clinical datum for presence or absence in the patient, by virtually testing with the algorithm both possible outcomes.

The *qualitative component* of benefit cannot be quantified; it depends on multiple factors such as patient health status and ability to tolerate the procedure, patient financial status, insurance company approval, prognosis, involved physician liability, and existence of efficacious and available treatments for the diseases listed in the differential diagnosis. Benefit must equal or exceed cost. The evaluation of cost-benefit of a clinical datum and the decision to implement a procedure to obtain it must be discussed with and approved by the patient. If the patient is wealthy, is not discouraged by the risk, or can tolerate discomfort, a procedure that incurs a greater cost may be acceptable. Confirmation of an uncertain diagnosis of a potentially life-threatening but treatable disease also may justify implementation of a more costly procedure.

Because cost and benefit cannot be accurately quantified, neither can cost-benefit ratio. If all of the aforementioned qualitative factors could be given an empiric value, it might be possible to devise an algorithm to assist the physician in better evaluating cost-benefit.

Summarizing how the three indices of clinical data—S, PP value, and cost—are determined: S depends on the clinical datum and corresponding disease; it is determined statistically with equation 2. PP value depends on S of the clinical datum for the disease under consideration and S of the same clinical datum for all diseases that manifest this datum; it is determined with equation 5. Cost—expense, risk, and discomfort—depends on the nature of the test or procedure needed to obtain the clinical datum; it is assigned one of four empirical categories: none, low, intermediate, or great.

Identifiers

Identifiers are assigned to a clinical datum when potentially dangerous circumstances prevail; they typically prompt a more extensive diagnostic work-up. An **interaction identifier** flags a clinical datum that might have been modified or masked by a drug interaction or a concurrent disease. The algorithm checks for such interactions. The use of this identifier is based on S of the clinical datum and will be explained later. A **risk identifier** flags any high-risk clinical datum (*e.g.*, dyspnea, chest pain, bleeding, laboratory test "panic" values) and any diagnosis with potentially bad outcome (*e.g.*, myocardial infarction, pulmonary embolism, malignancy.) Its purpose is to preclude overlooking such diagnoses; because the processing of these risk identifiers is not directly related to the mini-max procedure and because our diagnostic algorithm has other mechanisms to include the risky diagnoses in the differential diagnosis list, they are not discussed in this publication. More details can be found in our mentioned book [6.]

Before explaining how our algorithm relates to sensitivity, positive predictive value, and cost, we must discuss the ruling in and ruling out of diagnoses.

RULING IN AND RULING OUT DIAGNOSES

A diagnosis is *ruled in* when it is included in the differential diagnosis; this occurs whenever a patient clinical datum matches a clinical datum in the respective disease model.

A diagnosis is *ruled out* when it is deleted from the differential diagnosis; this occurs whenever the probability of the diagnosis falls below an empirical threshold. Clinical data that reduce the probability of a diagnosis favor this deletion.

These statements imply that a diagnosis must be ruled in before it can be ruled out.

When a new patient, for whom no clinical data are known comes to our attention, we consider two methods of achieving a final diagnosis:

The first method begins with a blank page—not knowing the patient's disease—and then formulating a differential diagnosis based on clinical data manifested by the patient. These patient clinical data, when matched with disease model clinical data, gradually increment the number of possible diagnoses; this process is called *ruling in diagnoses*. The greater PP value of a clinical datum that is *present*, the more likely the corresponding diagnosis. For example, microhemagglutination for *Treponema pallidum* test (MHA-TP) is a clinical datum of great PP value for syphilis; accordingly, if positive, it rules in this

disease with great probability, because few other diseases manifest this clinical datum. A clinical datum that is present, with great PP value, strongly rules in the diagnosis, even if S is small, meaning that this clinical datum is not frequently found; but as it is already present in this case, S is irrelevant. For example, filarias present in a blood sample is a clinical datum with great PP value for filariasis, confirming this diagnosis, despite a small S. On the other hand, a clinical datum present, typically would not favor a diagnosis only because it has a great S, because it simply tells that this clinical datum is frequently manifested by the specified disease, but many other diseases also may manifest it (small PP value). For example, weight loss has a great S for hyperthyroidism, but a small PP value; therefore, to rule in hyperthyroidism, a clinical datum with a greater PP value, such as suppressed thyroid stimulating hormone (TSH) must be investigated. A clinical datum that is present, with small S, typically would not rule in a diagnosis, because it simply means that this clinical datum is rare for the disease, which is not a reason *per se* to rule in the disease; for example, diarrhea (small S and small PP value) for duodenal ulcer. Accordingly, *ruling in a diagnosis relies on clinical data that are present and the greater the PP value* the more it will support this diagnosis. *S is irrelevant* if the clinical datum is present.

The second method of achieving a final diagnosis begins with all known diseases, gradually decrementing the number of possible diagnoses as sequentially investigated clinical data are found absent. An unknown patient could have *any* disease. For example, when we notice that he is a male, we realize that he cannot have an ovarian cancer; because he is young, prostate cancer is unlikely, and so forth. This process is called *ruling out* potential *diagnoses*. To *rule out* a potential diagnosis, we rely on the *sensitivity* of *clinical data* that are *absent* in the patient. The greater the S of a clinical datum that is absent, the less likely the corresponding diagnosis, even if the PP value is great, because the clinical datum is absent. For example, microhemagglutination test for *Treponema pallidum* (MHA-TP) is a clinical datum of great S for syphilis; accordingly, if negative, it rules out this disease because it is positive in essentially all cases of syphilis (rare false negative tests). As mentioned in the previous paragraph, weight loss is a clinical datum with great S for severe hyperthyroidism, because it is manifested in all such cases; accordingly, if this clinical datum is absent, the disease is ruled out. A clinical datum that is absent, with small S, has little influence on the probability of the diagnosis, even if PP value is great; for example, filarias negative in blood (great PP value, but small S) for filariasis. Small sensitivity of an absent clinical datum does not rule out the corresponding diagnosis because it only means that the clinical datum is rare for the disease; absence of a rare clinical datum does not exclude a diagnosis. For example diarrhea with small S and small PP value for duodenal ulcer, if absent, does not rule out this diagnosis. Accordingly, *ruling out a diagnosis relies on clinical data that are absent and with great S*; *PP value is irrelevant* if the clinical datum is absent.

This second method to achieve a final diagnosis is impractical because to analyze the several thousand known diseases would require a prohibitively large number of absent clinical data to rule out all but one or a few diseases. In reality, ruling in or ruling out is applied according to whether the clinical datum under consideration is present or absent, respectively.

In summary: A clinical datum present rules in the corresponding diagnosis with strength proportional to its positive predictive value (PP value). A clinical datum absent rules out the corresponding diagnosis with strength proportional to its sensitivity (S).

Table 2 shows how PP value and S of a clinical datum affect the P and ruling in or ruling out of a diagnosis according to whether the datum is present or absent in the patient.

Clinical datum	PP value	S	P	Effect on Diagnosis	Example	Comments
Present	Great	Great	Increased	Strongly ruled in	MHA-TP test for syphilis	Ruling in a diagnosis relies on clinical data present, with great PP value; S is irrelevant
	Great	Small	Increased	Strongly ruled in	Filariae in blood for filariasis	
	Small	Great	Unchanged	Weakly ruled in	Weight loss for hyperthyroidism	
	Small	Small	Unchanged	Weakly ruled in	Diarrhea for duodenal ulcer	
Absent	Great	Great	Decreased	Strongly ruled out	MHA-TP test for syphilis	Ruling out a diagnosis relies on clinical data absent, with great S; PP value is irrelevant
	Great	Small	Slightly changed according to value of S	Weakly ruled out	Filariae in blood for filariasis	
	Small	Great	Decreased	Strongly ruled out	Weight loss for hyperthyroidism	
	Small	Small	Slightly changed according to value of S	Weakly ruled out	Diarrhea for duodenal ulcer	

TABLE 2. MHA-TP, microhemagglutination for *Treponema palladium*, a highly exclusive and sensitive test for syphilis; PP value, positive predictive value; S, sensitivity; P, probability of diagnosis.

Eight combinations are possible—clinical datum present with great PP value, clinical datum present with great S, clinical datum present with small PP value, clinical datum present with small S, clinical datum absent with great PP value, clinical datum absent with great S, clinical datum absent with small PP value, and clinical datum absent with small S. Of these eight combinations, only two are useful—clinical datum present with great PP value and clinical datum absent with great S—because only they can significantly change the P of the corresponding diagnosis; all other combinations are discarded.

OPERATION OF OUR DIAGNOSTIC ALGORITHM

I. INITIAL CLINICAL DATA COLLECTION

The diagnostic process begins with collection of initial clinical data gleaned from the patient's history, physical examination, and prior consultations. These *initial* clinical data, entered in the computer, are unrefined because we do not yet know their PP value nor S. Each of these values depend on the specific clinical datum collected, but also on the corresponding, but as yet not ruled in diagnosis. Initially, collection is focused primarily on clinical data present because only these can rule in diagnoses. At this early phase, clinical data processing is purely categorical because we have not yet applied any probabilistic calculations; the diagnostic process is called *ill structured* [7]. Only after potential diagnoses are selected can S and PP value of clinical data be determined, the P of each diagnosis be calculated, and a differential diagnosis list be created; the diagnostic process then is said to be *well structured*. Then, if warranted, diagnoses can be ruled out by processing clinical data absent.

II. SELECTING POTENTIAL DIAGNOSES

Following initial clinical data collection, the algorithm next must compare each clinical datum manifested by the patient with all clinical data listed in all disease models stored in the knowledge base, selecting those disease models that contain one or more matching clinical data. Such disease models represent potential diagnoses that will become the differential diagnosis list. This task involves important difficulties that are, in our opinion, a major reason why a satisfactory diagnostic algorithm has not yet been achieved. On one hand, a disease typically never manifests all clinical data listed in its disease model; also, the cost of obtaining some of these clinical data may be prohibitive. On the other hand, similar clinical data can be manifested by diverse diseases; in other words, most clinical data are not pathognomonic. The algorithm selects all disease models that include one or more clinical data manifested by the patient. After selection of potential diagnoses, the algorithm must establish whether they are *competing* for a single final diagnosis or whether they correspond to *concurrent* diseases. Later on we will explain how the mini-max procedure deals with this problem.

III. CLINICAL DATUM LISTS

Having selected the matching disease models that now represent potential diagnoses, the algorithm creates, for each clinical datum present, a list that has for heading this clinical datum and comprises all potential diagnoses able to manifest such clinical datum.

A CLINICAL DATUM LIST is a list of diagnoses (*e.g.*, bronchitis, asthma, lung cancer) that a single clinical datum (*e.g.*, cough) evokes in the mind of the physician, which is analogous to the matching of a single clinical datum with clinical data in disease models by the computer. Such diagnoses, as opposed to diseases, are in the mind of the physician; the patient is not afflicted by all of them. Each of these *potential diagnoses* has a probability to become a *final diagnosis*, the latter ideally being concordant with the disease that afflicts the patient. A proper denomination for clinical datum list would be **Potential Diagnoses List for a Single Clinical Datum**, which is lengthy and cumbersome. For that reason, we abbreviate it to **Clinical Datum List**, because despite being a list of diagnoses, the clinical datum identifies the list and is its heading. Clinical Dat**um** List should not be mistaken with Disease Model, which has a disease as its heading and lists **all** clinical data that this disease can manifest.

Clinical Datum
Matching disease model → Potential diagnosis 1
Matching disease model → Potential diagnosis 2
⋮ ⋮
Matching disease model → Potential diagnosis n

The sensitivity (S) and positive predictive value (PP value) of the clinical datum for each potential diagnosis is shown. The diagnoses are sorted by decreasing PP value. In addition, an interaction identifier flags clinical data that may be modified by drugs or concurrent diseases and a risk identifier can flag clinical data and diagnoses that involve great risk or bad prognosis. Examples of clinical datum lists are:

Hemoptysis (bleeding from respiratory tract)	**S**	**PP value**
Tuberculosis	0.70	0.35 *
Lung cancer	0.40	0.20
Lung infarction	0.30	0.15
Bronchitis	0.05	0.025
Pneumonia	0.03	0.015
⋮	⋮	⋮

Dyspnea (difficulty to breathe)	**S**	**PP value**
Asthma	0.98	0.208
Congestive heart failure	0.80	0.17
Foreign body aspiration	0.80	0.17
Pneumonia	0.40	0.085
Emphysema	0.39	0.083
Carbon monoxide intoxication	0.22	0.047
Lung cancer	0.20	0.042
Lung infarction	0.19	0.04
Tuberculosis	0.17	0.036
Intense anemia	0.12	0.025
⋮	⋮	⋮

Cough	**S**	**PP value**
Foreign body aspiration	1.00	0.20
Bronchitis	0.98	0.196
Tuberculosis	0.70	0.14
Lung cancer	0.50	0.10
Lung infarction	0.40	0.08
Pneumonia	0.20	0.04
⋮	⋮	⋮

* The numeric values for S and PP value in the above examples were not obtained from actual statistics or calculations. Cost, risk and interaction identifiers were omitted for the sake of simplicity.

In a complete clinical datum list, the sum of the PP values equals 1. The fewer diagnoses a clinical datum list comprises, the more the clinical datum supports those diagnoses and the greater the corresponding PP values. When a clinical datum list contains only one diagnosis, PP value = 1, meaning the clinical datum is exclusive or pathognomonic for this diagnosis. Conversely, the more diagnoses a clinical datum

list comprises, the less the clinical datum supports those diagnoses, and the smaller their corresponding PP values.

IV. DIFFERENTIAL DIAGNOSIS LIST

This step creates a *differential diagnosis list* that includes *all diagnoses* listed in all clinical datum lists, without repeating similar diagnoses; in other words, all the potential diagnoses brought up thus far.

Here again, we have a semantic problem: a suitable name for DIFFERENTIAL DIAGNOSIS LIST could be **Potential Diagnoses List for More than one Clinical Datum** manifested by a patient. This denomination is long and cumbersome; we could abbreviate it to **Potential Diagnoses List**, but because DIFFERENTIAL DIAGNOSIS LIST is customary we retained the latter. Do not mistake Clinical Datum List with Differential Diagnosis List.

This method generates a relatively long differential diagnosis list and involves the burden of investigating the presence or absence of clinical data related to each diagnosis. However, this burden is not problematic because this method is applied at an early diagnostic stage when collected initial clinical data are already known to be present or absent, and the number of matched diagnoses is not excessive. Later, heuristically restrictive tools such as the best cost-benefit clinical datum next to investigate will further limit proliferation of clinical data and diagnoses. This method renders unnecessary the risk identifiers because it includes automatically both risk and non-risk diagnoses in the differential diagnosis list.

Example of differential diagnosis list based on the previous clinical datum lists:

Differential diagnosis list	**P**
Tuberculosis	0.35
Asthma	0.208
Lung cancer	0.20
Foreign body aspiration	0.20
Bronchitis	0.196
Congestive heart failure	0.17
Lung infarction	0.15
Pneumonia	0.085
Emphysema	0.083
Carbon monoxide intoxication	0.047
⋮	⋮

V. PROBABILITY OF DIAGNOSES. MINI-MAX PROCEDURE

At this point we have a well-structured diagnostic problem with a differential diagnosis list. Next, the algorithm must determine which of these potential diagnoses will become one or more final diagnoses.

Tomás Feder devised a procedure for calculating the probability (P) of a diagnosis by combining the PP value of clinical data when *present* (favoring a diagnosis) with the S of clinical data when *absent* (disfavoring a diagnosis). We call it *mini-max procedure**. In successive steps, we will explain this procedure with examples.

* Our term mini-max is reminiscent of a similar term used in game theory, but not previously applied in combination with Bayes formula to calculate probability of diagnoses (more details about this combination in our book [6].)

Step 1. Process clinical data present

To establish the value of P, other diagnostic programs add, subtract, multiply, or average the sensitivities, specificities, predictive values, estimated values of clinical data supporting a diagnosis, or iterate Bayes formula with each additional clinical datum. These approaches have flaws.

For example, jaundice, dark urine, and increased direct serum bilirubin are clinical data related by similar pathophysiologic mechanisms generated by a single lesion: biliary tract obstruction. Were we arithmetically to combine the individual PP values of these three equivalent and "redundant" clinical data, P of diagnosis biliary tract obstruction would be improperly increased thereby providing an undue advantage to this diagnosis, as compared to competing diagnoses. Furthermore, assume that an endoscopic retrograde cholangiopancreatography (ERCP) shows a biliary stone obstructing the common bile duct —a clinical datum that alone has a confirmatory PP value of 1. If we add the PP values of other supporting clinical data, the P of the diagnosis obstructing gallstones would exceed 1, which is probabilistically impossible. If we average or multiply these PP values, the confirmatory PP value 1 would be unduly reduced.

With our algorithm, the PP value of gallstone obstruction, which equals 1, supersedes all other clinical data with smaller PP value (jaundice, dark urine, increased serum bilirubin) because—whether present or absent—they would not change the diagnosis of obstructing gallstones already confirmed by ERCP. Accordingly, we consider that the greatest PP value of these related clinical data better represents the P of the diagnosis than any arithmetical combination of the individual values.

For each diagnosis in the differential diagnosis list, the algorithm looks in the entire set of clinical datum lists and selects the greatest PP value that supports *this* diagnosis. The selected greatest PP value equals the P of this diagnosis.

$$P_i = \max (\text{PP value}_1 \ldots \text{PP value}_i \ldots \text{PP value}_n) \quad (6)$$

Where P_i = probability of the diagnosis under consideration

max = maximum of

$\text{PP value}_1 \ldots \text{PP value}_i \ldots \text{PP value}_n$ = positive predictive values of clinical data present, that support the diagnosis under consideration

The algorithm then iterates the same routine to determine the P of each diagnosis in the differential diagnosis list.

Example: a patient presents with cough, hemoptysis, dyspnea, expectoration, and *Mycobacterium tuberculosis* (*Mycobacterium TB*) in sputum. Five clinical datum lists are generated:

Cough	**S**	**PP value**	**Hemoptysis**	**S**	**PP value**
Pulmonary tuberculosis	0.80	0.276	Pulmonary tuberculosis	0.40	0.222
Pulmonary embolism	0.50	0.172	Pulmonary embolism	0.60	0.333
Bronchiectasis	0.90	0.310	Bronchiectasis	0.30	0.167
Lung cancer	0.70	0.241	Lung cancer	0.50	0.278

Dyspnea	S	PP value
Pulmonary tuberculosis	0.20	0.148
Pulmonary embolism	0.50	0.370
Bronchiectasis	0.05	0.037
Lung cancer	0.60	0.444

Expectoration	S	PP value
Pulmonary tuberculosis	0.80	0.417
Pulmonary embolism	0.02	0.010
Bronchiectasis	0.90	0.469
Lung cancer	0.20	0.104

Mycobacterium TB	S	PP value
Pulmonary tuberculosis	0.70	1.000
Pulmonary embolism	0.00	0.000
Bronchiectasis	0.00	0.000
Lung cancer	0.00	0.000

S values in the above example are for demonstration purposes only and do not represent actual statistics. PP values were calculated by applying equation 5 to these S values. We assume that only the four listed diagnoses exist and that any of them could account for the five clinical data. Highlighted values refer to clinical data that are not elements of a specific disease model; accordingly, their S and PP value values equal 0. Such clinical data have no influence on calculated probabilities.

In the entire set of clinical datum lists, the greatest PP value for pulmonary tuberculosis is 1.000 and equals P for this diagnosis. Similarly, 0.370 for pulmonary embolism, 0.469 for brochiectasis, and 0.444 for lung cancer.

A differential diagnosis list is created, with each diagnosis showing the respective P that competes with the P of the other diagnoses for a final diagnosis. The diagnoses are sorted by decreasing P values.

Differential Diagnosis List	P
Pulmonary tuberculosis	1.000
Bronchiectasis	0.469
Lung cancer	0.444
Pulmonary embolism	0.370

Except for confirmed pulmonary tuberculosis, these P values do not yet satisfy thresholds that enable to rule out the other diagnoses or confirm some as concurrent final diagnosis (how such thresholds are determined will be explained later.) To satisfy this threshold requirement, the algorithm automatically determines which additional best cost-benefit clinical datum should next be investigated for its presence or absence. Fever is first recommended, followed by pulmonary cavity lesion:

Fever	S	PP value
Pulmonary tuberculosis	0.70	0.636
Pulmonary embolism	0.30	0.273
Bronchiectasis	0.00	0.000
Lung cancer	0.10	0.091

Cavity	S	PP value
Pulmonary tuberculosis	0.60	0.600
Pulmonary embolism	0.00	0.000
Bronchiectasis	0.10	0.100
Lung cancer	0.30	0.300

Were fever and a pulmonary cavity lesion also *present* in the patient, we would now have a total of seven clinical datum lists. Were the PP value associated with any of the diagnoses in these 2 new clinical datum lists to exceed the P of the same diagnosis, that greater PP value would replace this existing P.

PP value of a clinical datum present can only **increase** the probability of a diagnosis (equation 6.)

Step 2. Process clinical data absent

Only those clinical data absent that are related to diagnoses in the differential diagnosis list are processed.

We presented a rational explanation and example of why we believe that the greatest PP value of all the clinical data *present* that supports a specific diagnosis equals the P of this diagnosis. This is consequent to the fact that clinical data present are related by a common lesion or cause; adding the PP value of these clinical data excessively increases this P. We discussed how the sensitivity (S) of a clinical datum *absent* typically reduces the P of the corresponding diagnosis. To reduce P of the corresponding diagnosis, some authors arithmetically combine S of all absent clinical data, or sequentially apply Bayes formula to each S of such clinical data. We observed that this procedure excessively decreases the P of the diagnosis to a value that might incorrectly rule out the corresponding disease. For this reason, to reduce the P of a diagnosis, we use only the greatest S of all clinical data absent. This approach for disfavoring a diagnosis is less intuitive than using the greatest PP value of clinical data present for supporting a diagnosis. Clinical data absent are not related by a common lesion or cause; however, they might be related by a specific characteristic of patient's body that is responsible for the failure to react, or the cause is insufficient to evoke all potential clinical data. This common denominator justifies considering only the datum absent of greatest S as the representative of all clinical data absent. The next example supports this approach:

Consider again a patient with a suspected common bile duct obstruction by gallstones. An endoscopic retrograde cholangiopancreatography (ERCP) in this case was negative—*i.e.*, *no* gallstones were present in the common bile duct—an absent clinical datum of great S (close to 1) for the mentioned diagnosis. To rule out this diagnosis, it is unnecessary to consider additional clinical data absent of smaller S, such as right upper abdominal pain or vomiting.

So far, we have explained how clinical data present and their associated PP values determine the P of a diagnosis. Now we will explain how clinical data absent and their associated S values further influence this P.

Originally, we tried this equation:

$$P = \text{PP value} \times (1-S) \tag{7}$$

Where PP value = probability of a diagnosis **before** considering the sensitivity of a clinical datum absent; this probability equals the greatest PP value of all clinical data present that support this diagnosis (equation 6)

P = probability of the this diagnosis **after** considering the S of a clinical datum absent pertaining to the same diagnosis

S = sensitivity of a clinical datum absent pertaining to the same diagnosis

With equation 7, the greater the S of a clinical datum absent, the more it reduces the P of a diagnosis.

Let's assume that fever and cavity in a previous example, were investigated and found *absent* and let's apply equation 7. *Mycobacterium tuberculosis* was found in the sputum, a clinical datum present with a PP value = 1, which confers a P = 1 to tuberculosis, confirming this diagnosis. Next, our example considers fever, a clinical datum absent with S = 0.7. Applying equation 7, we obtain:

$$P = \text{PP value} \times (1-S) = 1 \times (1-0.7) = 0.3$$

Notice that the absence of fever decreases tuberculosis P from 1 to 0.3. It is unacceptable that a relatively unimportant clinical datum such as fever should cause a substantial decrease in P, which tends to rule out the already confirmed diagnosis of tuberculosis.

To temper this unacceptable decrease in P caused by equation 7, we instead use:

$$P_i = \frac{\text{PP value}_i\,(1-S_i)}{\text{PP value}_1\,(1-S_1) + \ldots + \text{PP value}_i\,(1-S_i) + \ldots + \text{PP value}_n\,(1-S_n)} \qquad (8)$$

Where P_i = probability of a diagnosis (*e.g.*, tuberculosis)

PP value_i = positive predictive value of the clinical datum present (*e.g.*, *Mycobacterium tuberculosis* in sputum for tuberculosis)

S_i = sensitivity of the clinical datum absent (fever for tuberculosis)

$\text{PP value}_1 \ldots \text{PP value}_i \ldots \text{PP value}_n$ = positive predictive value of the same clinical datum present (*Mycobacterium tuberculosis* in sputum) for each respective diagnosis in the differential diagnosis list (4 diagnoses per our example)

$S_1 \ldots S_i \ldots S_n$ = sensitivity of the clinical datum absent (fever) for each respective diagnosis in the differential diagnosis list (4 diagnoses per our example)

Notice that the numerator of equation 8 is identical to the right member of equation 7, and that a denominator has been introduced, the effect of which is to "temper" the result. This denominator comprises several terms, each of which refers to a diagnosis in the differential diagnosis list. Each comprises the PP value of the clinical datum present (*Mycobacterium tuberculosis*) and the S of the clinical datum absent (fever). These clinical data present and absent remain *unchanged for all terms*; but their respective PP values and S values change to values associated with each diagnosis. Equation 7 is then applied to these values in each denominator term of equation 8.

Equation 8 is related to Bayes formula, but is here used differently than in other programs; it involves two clinical data supposedly *independent*—one present and the other absent. Accordingly, Bayes condition of independence is not violated.

Referring to our example of *Mycobacterium tuberculosis* present in sputum and fever absent, we now apply equation 8 to calculate the probability of tuberculosis (P_{TB}):

$$P_{TB} = \frac{\text{PP value}_{TB}\,(1-S_{TB})}{\text{PPvalue}_{TB}(1-S_{TB})+\text{PPvalue}_{\text{bronchiectasis}}(1-S_{\text{bronchiectasis}})+\text{PPvalue}_{\text{cancer}}(1-S_{\text{cancer}})+\text{PPvalue}_{\text{embolism}}(1-S_{\text{embolism}})}$$

Substituting PP value and S with values from the clinical datum lists, we obtain:

$$P_{TB} = \frac{1.00\,(1-0.70)}{1.00\,(1-0.70) + 0.00\,(1-0.00) + 0.00\,(1-0.10) + 0.00\,(1-0.30)} =$$

$$= \frac{0.30}{0.30 + 0.00 + 0.00 + 0.00} = 1.00$$

Notice that equation 8 retains the correct value of P = 1 for confirmed tuberculosis, instead of P = 0.30, as was obtained with equation 7.

Equation 8 yields the same result if all PP values of the clinical datum *present* are substituted with the corresponding S′ of the same clinical datum, S otherwise typically used with clinical data *absent*:

$$P_i = \frac{\text{PP value}_i\,(1-S_i)}{\text{PP value}_1\,(1-S_1) + \ldots + \text{PP value}_i\,(1-S_i) + \ldots + \text{PP value}_n\,(1-S_n)} = \frac{S_i'\,(1-S_i)}{S_1'\,(1-S_1) + \ldots + S_i'\,(1-S_i) + \ldots + S_n'\,(1-S_n)}$$

Notice that the value of S′ (sensitivity of the clinical datum *present*) is *not* the same as the value of S (sensitivity of the clinical datum *absent*). Equation 8, with PP values, yields the same result as with S′ because PP value and S′ of a given clinical datum for a given diagnosis are directly proportional. When all PP values are substituted with the right member of equation 5, equation 8 can be simplified to its substituted form (right member) shown above. This simplification is possible because the sum $S_1+\ldots+S_i+\ldots+S_n$ in the numerator and denominator of equation 8 have the same values and cancel each other.

The denominator of equation 8—modified Bayes formula—can be seen as a *weighted average* of S′ of a clinical datum present for diverse diseases, weighted by S values of the clinical datum absent. Comparing the S′ of the clinical datum present for a specific disease with the average of S′ values for all diseases indicates the relative significance of the clinical datum for this specific disease.

For clarity and consistency, we will retain the original equation 8 (with PP values) for all further calculations, but the substituted form (with S′ values) might be useful for computer programming.

Equation 8 is then iterated to calculate the P that the clinical data pair *Mycobacterium tuberculosis*-fever confers to the remaining diagnoses in the differential diagnosis list. PP value and S value corresponding to each diagnosis must be substituted in the numerator; the denominator remains unchanged. Equation 8 *normalizes* the probabilities of the diagnoses, meaning that their sum ($P_{TB} + P_{bronchiectasis} + P_{cancer} + P_{embolism}$) now equals 1. Referring to our example:

***Mycobacterium* TB - Fever**	**PP value**	**S**		**P**
Pulmonary tuberculosis	1.000 × (1-0.70) =	0.300 (numerator)	÷ 0.300 (denominator) =	1.000
Pulmonary embolism	0.000 × (1-0.30) =	0.000 (numerator)	÷ 0.300 (denominator) =	0.000
Bronchiectasis	0.000 × (1-0.00) =	0.000 (numerator)	÷ 0.300 (denominator) =	0.000
Lung cancer	0.000 × (1-0.10) =	0.000 (numerator)	÷ 0.300 (denominator) =	0.000
		Sum = 0.300 (denominator)		Sum = 1.000

Step 3. Create clinical data pairs

As remarked above, equation 8 comprises a clinical datum present and a clinical datum absent; we call this clinical data combination a *clinical data pair*. Each clinical data pair confers a *partial probability* to a diagnosis. To calculate the *total probability* of each diagnosis, the mini-max procedure must create *all possible clinical data pairs* with all thus-far investigated clinical data present and absent. The number of clinical data pairs created will equal the number of clinical data present multiplied by the number of clinical data absent.

Returning to our previous example, we had 5 clinical data present (cough, expectoration, hemoptysis, dyspnea, and *Mycobacterium tuberculosis*) and 2 clinical data absent (cavity and fever), creating a total of 10 clinical data pairs (cough-cavity, cough-fever, hemoptysis-cavity, hemoptysis-fever, dyspnea-cavity, dyspnea-fever, expectoration-cavity, expectoration-fever, *Mycobacterium tuberculosis*-cavity, and *Mycobacterium tuberculosis*-fever.)

Because PP value and S value vary with each diagnosis, the total number of resulting *partial P values* equal the number of clinical data pairs created multiplied by the number of diagnoses in the differential diagnosis list. In our example, we had 10 clinical data pairs and 4 diagnoses (pulmonary tuberculosis, pulmonary embolism, bronchiectasis, and lung cancer), yielding a total of 40 partial P values.

Step 4. Create clinical data pair tables

We then organize the 40 partial P values as 10 *clinical data pair tables*, one table for each clinical data pair. Each table is headed by the clinical data pair; its first column lists the diagnoses of the differential diagnosis list; intermediate columns apply equation 8, and its last column lists the resultant partial P values. For our example:

Clinical data pair tables

Cough-Cavity	**PP value**		**S**			**Partial P**
Pulmonary tuberculosis	0.276	× (1−0.60) =	0.110	÷ 0.730 =		0.151
Pulmonary embolism	0.172	× (1-0.00) =	0.172	÷ 0.730 =		0.236
Brochiectasis	0.310	× (1-0.10) =	0.279	÷ 0.730 =		0.382
Lung cancer	0.241	× (1-0.30) =	0.169	÷ 0.730 =		0.231
			0.730			1.000
Cough-Fever						
Pulmonary tuberculosis	0.276	× (1-0.70) =	0.083	÷ 0.731 =		0.113
Pulmonary embolism	0.172	× (1-0.30) =	0.121	÷ 0.731 =		0.165
Brochiectasis	0.310	× (1-0.00) =	0.310	÷ 0.731 =		0.425
Lung cancer	0.241	× (1-0.10) =	0.217	÷ 0.731 =		0.297
			0.731			1.000
Hemoptysis-Cavity						
Pulmonary tuberculosis	0.222	× (1-0.60) =	0.089	÷ 0.767 =		0.116
Pulmonary embolism	0.333	× (1-0.00) =	0.333	÷ 0.767 =		0.435
Brochiectasis	0.167	× (1-0.10) =	0.150	÷ 0.767 =		0.196
Lung cancer	0.278	× (1-0.30) =	0.194	÷ 0.767 =		0.254
			0.767			1.000
Hemoptysis-Fever						
Pulmonary tuberculosis	0.222	× (1-0.70) =	0.067	÷ 0.717 =		0.093
Pulmonary embolism	0.333	× (1-0.30) =	0.233	÷ 0.717 =		0.325
Brochiectasis	0.167	× (1-0.00) =	0.167	÷ 0.717 =		0.233
Lung cancer	0.278	× (1-0.10) =	0.250	÷ 0.717 =		0.349
			0.717			1.000

Dyspnea-Cavity

Pulmonary tuberculosis	0.148	×	(1-0.60)	=	0.059	÷	0.774	= 0.077
Pulmonary embolism	0.370	×	(1-0.00)	=	0.370	÷	0.774	= 0.478
Brochiectasis	0.037	×	(1-0.10)	=	0.033	÷	0.774	= 0.043
Lung cancer	0.444	×	(1-0.30)	=	0.311	÷	0.774	= 0.402
					0.774			1.000

Dyspnea-Fever

Pulmonary tuberculosis	0.148	×	(1-0.70)	=	0.044	÷	0.741	= 0.060
Pulmonary embolism	0.370	×	(1-0.30)	=	0.259	÷	0.741	= 0.350
Bronchiectasis	0.037	×	(1-0.00)	=	0.037	÷	0.741	= 0.050
Lung cancer	0.444	×	(1-0.10)	=	0.400	÷	0.741	= 0.540
					0.741			1.000

Expectoration-Cavity

Pulmonary tuberculosis	0.417	×	(1-0.60)	=	0.167	÷	0.672	= 0.248
Pulmonary embolism	0.010	×	(1-0.00)	=	0.010	÷	0.672	= 0.016
Bronchiectasis	0.469	×	(1-0.10)	=	0.422	÷	0.672	= 0.628
Lung cancer	0.104	×	(1-0.30)	=	0.073	÷	0.672	= 0.109
					0.672			1.000

Expectoration-Fever

Pulmonary tuberculosis	0.417	×	(1-0.70)	=	0.125	÷	0.695	= 0.180
Pulmonary embolism	0.010	×	(1-0.30)	=	0.007	÷	0.695	= 0.010
Bronchiectasis	0.469	×	(1-0.00)	=	0.469	÷	0.695	= 0.675
Lung cancer	0.104	×	(1-0.10)	=	0.094	÷	0.695	= 0.135
					0.695			1.000

***Mycobacterium TB* -Cavity**

Pulmonary tuberculosis	1.000	×	(1-0.60)	=	0.400	÷	0.400	= 1.000
Pulmonary embolism	0.000	×	(1-0.00)	=	0.000	÷	0.400	= 0.000
Bronchiectasis	0.000	×	(1-0.10)	=	0.000	÷	0.400	= 0.000
Lung cancer	0.000	×	(1-0.30)	=	0.000	÷	0.400	= 0.000
					0.400			1.000

***Mycobacterium TB* -Fever**

Pulmonary tuberculosis	1.000	×	(1-0.70)	=	0.300	÷	0.300	= 1.000
Pulmonary embolism	0.000	×	(1-0.30)	=	0.000	÷	0.300	= 0.000
Bronchiectasis	0.000	×	(1-0.00)	=	0.000	÷	0.300	= 0.000
Lung cancer	0.000	×	(1-0.10)	=	0.000	÷	0.300	= 0.000
					0.300			1.000

Step 5. Calculate partial P that each clinical data pair confers to each diagnosis

To calculate the partial P that each clinical data pair confers to each diagnosis in the differential diagnosis list, equation 8 is applied to the PP value of the clinical datum present and the S of the clinical datum absent for each diagnosis (see clinical data pair tables above.)

Notice that throughout these iterations of equation 8, the entire denominator remains unchanged for each clinical data pair and table. However, the numerator does change with each iteration; it assumes the value of the denominator term corresponding to the diagnosis being processed.

Step 6. Create mini-max tables

Now, we must determine the *total probability* that the partial probabilities mentioned in steps 3, 4 and 5 confer to each diagnosis in the differential diagnosis list. This is achieved by creating a *mini-max table* for each diagnosis (see below.)

The first column of each mini-max table lists each clinical datum present. The second column lists the PP value of each clinical datum present; its bottom cell repeats the greatest of these values, which is the total P of the diagnosis *before* clinical data absent are considered. The next several columns show the partial P values that each clinical data pair confers to the diagnosis; the number of these columns equals the number of clinical data absent. The heading of each column shows the clinical datum absent and its S for the diagnosis. Each partial P value is transferred from the clinical data pair table to the mini-max table cell where the clinical data present and absent converge. The bottom cell of each column repeats the greatest partial P value appearing in the column. The last column repeats the smallest value appearing in each row. The bottom cell of this column, which also is the last cell of the mini-max diagnosis table, repeats the greatest value of the column; it equals the total P of the diagnosis, *after* clinical data absent have been considered.

Mini-max tables

Mini-max table for tuberculosis

TUBERCULOSIS	PP value = partial P before considering clinical data absent	Partial P with Cavity absent S = 0.6	Partial P with Fever absent S = 0.7	MINIMUM VALUE IN EACH ROW
Cough present	0.276	0.151	0.113	0.113
Hemoptysis present	0.222	0.116	0.093	0.093
Dyspnea present	0.148	0.077	0.060	0.060
Expectoration present	0.417	0.248	0.180	0.180
MTb present	1.000	1.000	*1.000*	1.000
MAXIMUM VALUE IN EACH COLUMN	1.000	1.000	1.000	**Total P = 1.000**

MTb, *Mycobacterium tuberculosis*. The total probability of **tuberculosis** at this diagnostic step is the maximum value (**1.000**) in the last column.

Mini-max table for pulmonary embolism

PULMONARY EMBOLISM	PP value = partial P before considering clinical data absent	Partial P with Cavity absent S = 0	Partial P with Fever absent S = 0.3	MINIMUM VALUE IN EACH ROW
Cough present	0.172	0.236	0.165	0.165
Hemoptysis present	0.333	0.435	0.325	0.325
Dyspnea present	0.370	0.478	*0.350*	0.350
Expectoration present	0.010	0.016	0.010	0.010
MTb present	0.000	0.000	0.000	0.000
MAXIMUM VALUE IN EACH COLUMN	0.370	0.478	0.350	**Total P = 0.350**

The total probability of **pulmonary embolism** at this diagnostic step is the maximum value (**0.350**) in the last column.

Mini-max table for bronchiectasis

BRONCHIECTASIS	PP value = partial P before considering clinical data absent	Partial P with Cavity absent S = 0.1	Partial P with Fever absent S = 0	MINIMUM VALUE IN EACH ROW
Cough present	0.310	0.382	0.425	0.310
Hemoptysis present	0.167	0.196	0.233	0.167
Dyspnea present	0.037	0.043	0.050	0.037
Expectoration present	*0.469*	*0.628*	0.675	0.469
MTb present	0.000	0.000	0.000	0.000
MAXIMUM VALUE IN EACH COLUMN	0.469	0.628	0.675	**Total P = 0.469**

The total probability of **bronchiectasis** at this diagnostic step is the maximum value (**0.469**) in the last column. Were the second column not included in the calculation, total P of this diagnosis would be 0.628; see property 4 B of mini-max procedure (explained later.)

Mini-max table for lung cancer

LUNG CANCER	PP value = partial P before considering clinical data absent	Partial P with Cavity absent S = 0.3	Partial P with Fever absent S = 0.1	MINIMUM VALUE IN EACH ROW
Cough present	0.241	0.231	0.297	0.231
Hemoptysis present	0.278	0.254	0.349	0.254
Dyspnea present	0.444	*0.402*	0.540	0.402
Expectoration present	0.104	0.109	0.135	0.104
MTb present	0.000	0.000	0.000	0.000
MAXIMUM VALUE IN EACH COLUMN	0.444	0.402	0.540	**Total P = 0.402**

The total probability of **lung cancer** at this diagnostic step is the maximum value (**0.402**) in the last column.

Step 7. Determine total P of a diagnosis

In the mini-max table, the last column lists the smallest values of each row; the greatest value in this last column, repeated in the last cell of the table, equals the total P of the diagnosis. Therefore, the algorithm determines the total P of a diagnosis based on partial P values; it involves the following concepts:

1. A clinical data pair comprises a clinical datum present and a clinical datum absent.
2. A specific clinical data pair, that we call *determining clinical data pair*, determines the total P of a specific diagnosis.
3. In this determining clinical data pair, applying equation 8 to the PP value of the clinical datum present and the S of the clinical datum absent, yields a specific partial P that we call *determining partial P* because it determines and equals the total P of the specific diagnosis.
4. A *specific cell* for this *determining partial P* exists in the mini-max table of this specific diagnosis.
    - In this specific cell, the mentioned clinical datum present converges with the mentioned clinical datum absent.
    - In this specific cell, the value (italicized) of the determining partial P was transferred from its clinical data pair table to the corresponding mini-max table.
    - In this specific cell, the value of the determining partial P is at once the smallest in its row and the greatest in its column (see below.)

To find the clinical data pair (determining clinical data pair) responsible for the current total P of a diagnosis, we must backtrack the steps that led from that pair to the total P. Start at the last cell (total P) of the mini-max table and ascend (following the arrows) to any cell with the same value, then go left on that row until any cell with the same value (the determining partial P) is encountered. The clinical datum present and the clinical datum absent that converge to this cell comprise the requisite clinical data pair (determining clinical data pair); the respective PP value and S are shown in the mini-max table.

Example: In the mini-max table of Lung Cancer the last cell shows the current total P (0.402) of this diagnosis. Following the arrows takes us to another cell with the italicized value 0.402, which is the determining partial P. To this cell converge PP value (0.444) of dyspnea present and S (0.3) of pulmonary cavity absent. The determining clinical data pair dyspnea-cavity is responsible for the current determining partial P and total P (0.402.)

Mini-max tables are not based on Bayes formula and therefore circumvent the problem of clinical data independence and disease incompatibility.

Alternative methods exist for determining the total P of a diagnosis:

- *Mini-max* alternative establishes the *maximum partial P* in the last column as the *total P of the diagnosis*.
- *Maxi-min* alternative establishes the *minimum partial P* value in the last row as the *total P of the diagnosis*.

In all cases thus far examined, these two alternatives have yielded identical total P values; but playing with bizarre fictitious sensitivity values, we fabricated an example where the mini-max and the maxi-min alternatives yielded different total P values. It can be demonstrated that with only two diagnoses in the differential diagnosis list, mini-max and maxi-min alternatives always yield an identical result; when more than two diagnoses are processed, different results may be obtained on rare occasions. This demonstration is beyond the scope of this publication. Searching for a general principle that can predict when the two probabilities will be equal, Tomás Feder found three pertinent rules:

1. If a row exists in which each cell contains the maximum value in the corresponding column.

2. If a column exists in which each cell contains the minimum value in the corresponding row.
3. If and only if a cell exists that at once contains the minimum value in its row and the maximum value in its column; this value (determining partial P) equals the value of the diagnosis total probability.

If only one rule is satisfied, total P is equal for both mini-max and maxi-min alternatives. Rule 1 and rule 2 are sufficient but not necessary; they are too strong and imply rule 3. Rule 3 is necessary and sufficient.

Unless experimentation with a prototype program proves otherwise, actual clinical cases almost always comply with rule 3, yielding the same total P with both alternatives. We ascribe this behavior to the typical proportion among sensitivities and the fact that the partial P values in the rows of the mini-max table present a monotone relation. The latter means that when in one row the partial P value increases or decreases from one cell to the next, in the other rows the changes occur in the same direction.

Were such monotone relation broken and were the mini-max and maxi-min alternatives to yield different results, a possibility could exist that a clinical datum *present* with a PP value *smaller* than the current greatest PP value of clinical data present could still *increase* the total P of the diagnosis. This would contradict the general definition of the P of diagnosis (equation 6.)

If monotony is broken, a determining partial P that is at the same time the greatest partial P of its column and the smallest partial P of its row may not exist. In this case, two different partial P will exist, one determining the total P of the diagnosis with the mini-max alternative and another determining the total P with the maxi-min alternative. Were an incidental clinical case to occur, where mini-max and maxi-min results differ, we would favor the first alternative. The last column of the mini-max table lists partial P values that remain from *clinical data present* after being reduced by clinical data absent, whereas the last row lists partial P values of *clinical data absent* after being increased by clinical data present. It seems more reasonable to equal total P to the greatest partial P of clinical data present (mini-max alternative) rather than to the smallest partial P of clinical data absent (maxi-min alternative.)

- Average of mini-max and maxi-min alternative results (if unequal) as the *total P of the diagnosis*.

- Nash equilibrium result as the *total P of the diagnosis*: The mini-max procedure can be seen as a two-player game of antagonist clinical data present versus clinical data absent, where a novel combination of Nash equilibrium and a modified Bayes formula find its application to medical diagnosis. In the mini-max table, a maximizing player chooses a row for a datum present and tries to maximize the resulting value, and a minimizing player chooses a column for a datum absent and tries to minimize the resulting value. The cell where the chosen row and column intersect shows the resulting value of equation 8 applied to the corresponding PP value and S. In the mini-max alternative, the maximizing player is first in choosing the row and then the minimizing player chooses the column. In the maxi-min alternative, the minimizing player is first in choosing the column and then the maximizing player chooses the row. In both cases, the second player, knowing the selection of the first, is able to counter it with the optimal strategy (this is a deterministic strategy called pure strategy). The strategy becomes uncertain if both players must choose their rows and columns simultaneously, because none of them knows the strategy chosen by the other. In such a situation the maximizing player may assign to the rows, probabilities adding to 1, and the minimizing player may assign to the columns, probabilities adding to 1 (this is a probabilistic strategy called mixed strategy.) If a row has assigned probability p and a column has assigned probability q, the *probability of the intersecting cell to be chosen* at random is $p \times q$, a value different from the *partial P*, result of equation 8 applied to the corresponding clinical data present and absent. Multiplying $p \times q \times$ partial P for each cell and adding the results for all cells in the mini-max table, yields the expected value of the chosen cell (in our case the total P of the diagnosis.) One can show that such a game always has a Nash equilibrium (P, Q), that is a choice of a set P of probabilities p for the set of rows and a set Q of probabilities q for the set of columns (as described above), so that when the maximizing player chooses P the resulting value is minimized by Q, and when the minimizing player chooses Q the resulting value is maximized by P. Nash equilibrium calculates the optimal strategy determining P ($p_1 \ldots p_n$) and Q ($q_1 \ldots q_n$). The Nash equilibrium (P, Q) yields a value always between mini-max and maxi-min results, but different from the average of both. When mini-max and maxi-min results are equal, which is the typical scenario, the game is strictly determined, and the expected value calculated with Nash equilibrium equals this result. Ideally, one should always determine total P of the diagnosis by computing the Nash equilibrium. However, determination of sets P and Q is mathematically complex and could invest a greater computing time, without offering a significant advantage over the simpler mini-max alternative, given the negligible or no difference between mini-max and maxi-min results. We presented the Nash equilibrium alternative only for the theoretical interest of combining it with Bayes formula, in a mini-max environment and as a stimulus for computer researchers to find other applications where probabilistic accuracy is more critical.

Step 8. Update the differential diagnosis list

Next, we again sort all diagnoses in the differential diagnosis list, according to decreasing total P values:

Differential Diagnosis List	Total P
Tuberculosis	1.000
Bronchiectasis	0.469
Lung cancer	0.402
Pulmonary embolism	0.350

Notice that the total P values of the diagnoses have changed and are more widely dispersed, but they still do not satisfy, except for tuberculosis, our threshold requirements for confirming as final or ruling out each of the other diagnoses. Accordingly, additional clinical data must be investigated. Then the mini-max procedure must be iterated with each additional clinical datum, and the total P of all diagnoses recalculated, until requirements for conclusion of the diagnostic quest are satisfied.

***

Properties of the mini-max procedure

1. Each additional clinical datum present generates a new *row* in an existing mini-max table.

2. Each additional clinical datum absent generates a new *column* in an existing mini-max table.

3. A mini-max table (in typical cases where mini-max and maxi-min results are equal) has only one determining partial P cell, the value of which (italicized in the table) is *the smallest of its row and the greatest of its column*. The clinical datum present and the clinical datum absent that converge to this cell constitute the determining clinical data pair that originated this determining partial P value that equals the *total P value* of the diagnosis.

4. When an additional clinical datum *present* is processed with the mini-max procedure, the total P of the diagnosis may increase, depending on its PP value. Typically, when an additional clinical datum *absent* is processed with the mini-max procedure, the greater its S, the more it decreases the P of the diagnosis. However, exceptions to this rule result from the effect this S has on the partial P of the other diagnoses that share the clinical data pair table, and from the interaction of the resulting partial P values in the mini-max table. The total P of the diagnosis will either decrease, increase, or remain unchanged:

   A. *Total P decreases*. Let's concentrate on a clinical data pair table. For a specific diagnosis, the S value of the clinical datum absent is *inversely* related to *its* partial P and *directly* related to the partial P values of the *other* diagnoses. An additional clinical datum absent typically reduces the total P of a diagnosis if its S is greater than the S of the absent clinical datum in the *determining clinical data pair*, in turn responsible for the current *determining partial P* of the diagnosis. If this condition is fulfilled, this new partial P will be smaller than the current determining partial P and becomes the new determining partial P that equals total P in the mini-max table.

   B. *Total P increases*. The mini-max procedure is not intended to increase the total P of a diagnosis based on clinical data absent. Nevertheless, this occasionally occurs, but only when a *first* clinical datum absent is processed; because at this point only one clinical datum absent column is

generated, smaller values do not exist in the rows. The greatest partial P value in this column becomes the determining partial P and if it exceeds the current total P, it will replace the latter. Any subsequent clinical datum absent that is processed—regardless of its S value and resulting partial P—can only decrease the total P, because only the smallest partial P in a row can become a determining partial P. If we do not want a first clinical datum absent to increase the current total P of a diagnosis, then the second column of the mini-max table must be included in the calculation. In this way, we avoid violating the general rule that a clinical datum absent must never increase the total P. However, it is questionable whether an occasional increase due to a clinical datum absent is illegitimate and should be disallowed.

An example is the mini-max table for bronchiectasis (see table above), where the total P of this diagnosis would have been 0.628 (italicized) instead of 0.469, were the second column not included in the calculation.

C. *Total P does not change*, when a clinical datum absent does not fulfil any of the conditions for decreasing or increasing total P. This occurs frequently; furthermore, the total P of a diagnosis is quite resistant to change, especially for diagnoses with a great total P. This is an important advantage of the mini-max procedure, because it precludes ruling out a confirmed diagnosis (strongly supported by clinical data present) by some relatively unimportant clinical datum absent (as seen above in the example of tuberculosis.)

5. The order in which clinical data are processed is irrelevant; it will change only the relative position of the generated new row or column without affecting the total probability of the diagnosis. This commutative order is intuitive and consistent with experience.

6. When an additional clinical datum present or absent is incorporated into a mini-max table, the previously calculated partial P values of the diagnosis in the table remain unchanged and need not be recalculated. Such P values are retained in case a need arises to determine which clinical datum pair generated a partial probability in a cell. The algorithm need remember the values in the last column only. Whenever an additional clinical datum is processed, new clinical data pairs are generated and new partial P values are calculated. The algorithm then compares these new partial P values with the existing partial P values in the last column and calculates the new total P of the diagnosis.

7. An interesting property of the mini-max procedure is revealed when the sum of the total P of all diagnoses in the differential diagnosis list is substantially greater than 1; it suggests that not all such diagnoses are competing, but that some represent concurrent diseases. The degree of support that a clinical datum gives to a diagnosis is directly proportional to its corresponding PP value. This value can be found in the clinical datum list associated with the diagnosis or in the second column of the mini-max table. If all clinical data predominantly support the same diagnosis, the remaining diagnoses tend to compete and the sum of the probabilities of all diagnoses in the differential diagnosis list is close to 1. When some clinical data predominantly favor one diagnosis and other clinical data predominantly favor another diagnosis, these diagnoses tend to be concurrent; the sum of their probabilities will be considerably greater than 1. Concurrent diagnoses are supported by different clinical data; accordingly, each concurrent diagnosis can by itself attain a probability up to 1. The greater the sum of the probabilities, the greater the number of concurrent diseases.

8. When a clinical datum present *with* a PP value that approaches or equals 1 strongly supports or confirms a diagnosis, a clinical datum absent—regardless of its S value—*cannot* reduce the great P that such a clinical datum present confers to the diagnosis. This property also is true for concurrent diagnoses with great probability in the differential diagnosis list. This important advantage precludes

a confirmed diagnosis from being ruled out by a relatively unimportant clinical datum absent (see tuberculosis example.) However, the P of a diagnosis *without* a confirming clinical datum present *may* be reduced by the S of such a clinical datum absent. Retaining diagnoses with great P, while simultaneously ruling out diagnoses with a small P, enables concurrent diagnoses to be distinguished from competing diagnoses. The manner in which the algorithm processes concurrent diagnoses will be addressed later.

9. Occasionally, when an additional clinical datum changes the total P value of a diagnosis in a differential diagnosis list, it may result in a significant opposite P change in one or more of the other diagnoses. Should such results be accepted or disregarded? Acceptance of the P changes in the other diagnoses contradicts the rule that each clinical datum should be processed for its corresponding diagnosis only. Experimentation with a prototype diagnostic program perhaps will solve this dilemma. In any event, the greater the total P of any diagnosis, the more it resists change.

10. Each time an additional clinical datum becomes available, the entire mini-max procedure is iterated recalculating total P values of the diagnoses, simultaneously processing all present and absent clinical data. However, unchanged values in mini-max table cells do not need to be recalculated (property 6.)

What happens when a diagnosis with a great P, based on a clinical datum present with a PP value = 1, is confronted with an additional clinical datum absent with an S = 1? Would the clinical datum present or the clinical datum absent win the rule in/rule out contest? In an actual case, this confrontation would be most unlikely to occur because S = 1 means that this clinical datum is always present, contradicting its absence. Furthermore, the mini-max procedure precludes the discredit of a diagnosis with a great total P by a clinical data absent (property 8 of mini-max procedure.)

A situation similar to that in the previous paragraph occurs when a diagnosis is evident, but an important clinical datum that is expected to be present is erroneously absent or unobserved. For example a patient with polyuria, polydipsia, and strong family history of diabetes mellitus does not manifest the expected hyperglycemia due to a laboratory error. Diabetes insipidus, psychogenic excessive water ingestion, renal disease, and other causes for the mentioned clinical data have been ruled out. Perhaps a routine should be included in the algorithm that detects such contradictions and recommends verification of the discordant clinical datum.

VI. BEST COST-BENEFIT CLINICAL DATUM NEXT TO INVESTIGATE

The best cost-benefit clinical datum next to investigate for presence or absence in a patient is an important function that can substantially shorten and reduce the cost of a diagnostic quest by precluding investigation of futile clinical data. This has important socioeconomic implications, especially in this era of managed care, when insurance companies curtail tests and procedures, and when physicians are rated by their proficiency in ordering tests in general.

Computers are faster and more accurate than the human brain in selecting the most convenient clinical datum next to investigate for presence or absence at each diagnostic step.

Our algorithm recommends the best cost-benefit clinical datum next to investigate, based on cost, PP value, and S.

Because the term *best cost-benefit clinical datum next to investigate in a patient* is lengthy, we shorten it to **best cost-benefit clinical datum**.

The best cost-benefit clinical datum function enables us to predict which new clinical datum will most increase or decrease the total probability (P) of a diagnosis, reducing the number of clinical data required to achieve a final diagnosis.

A recommended best cost-benefit clinical datum can be evaluated—before actually accomplishing the corresponding test or procedure—by *virtually* considering it either present or absent, while observing how much it improves the diagnostic outcome.

Initial clinical data collection was achieved during the history and physical examination. We accepted whatever clinical data were revealed, without considering their rule-in or rule-out power. Subsequent clinical data collection is more selective, because we have a differential diagnosis list and a better-structured diagnostic process that enables to apply statistical and probabilistic concepts, and choose the best cost-benefit clinical datum, based on cost, positive predictive value, and sensitivity.

To select a best cost-benefit clinical datum, several steps must be followed:

Step 1. Select clinical data not yet investigated in the patient

The algorithm examines *every* diagnosis in the differential diagnosis list and selects from its respective disease model all clinical data *not yet investigated*. These clinical data differ from those initially collected; they are expected to be numerous because each disease model will contribute many new clinical data. However, only clinical data of appropriate cost and either of great PP value or great S need be investigated for presence or absence.

Step 2. Organize clinical data not yet investigated according to cost category, diagnosis, PP value, and S

Clinical data not yet investigated are organized according to three hierarchical levels (Fig. 1, next page.)

The *first level* is **COST CATEGORIES**, comprising four categories: none, small, intermediate, and great. The *second level* is **DIAGNOSES**, comprising all diagnoses in the differential diagnosis list, identically repeated in each cost category, in order of decreasing P value. The *third level* is **CLINICAL DATA**, comprising two lists that we call **PP VALUE LIST** and **S LIST** containing only those clinical data that have a cost similar to the corresponding cost category. Both lists contain the same clinical data, but ordered according to decreasing PP value and decreasing S value respectively, and consequently these clinical data are shown with different sequence in each list.

Step 3. Recommend a new clinical datum as best clinical datum assuming it present

The algorithm moves to the lowest as-yet-unprocessed COST category, selects the as-yet-unprocessed DIAGNOSIS with greatest P, and from the corresponding PP VALUE LIST, selects the as-yet-unprocessed clinical datum with the greatest PP value. This PP value then is compared to the PP value of the clinical datum present in the *current* determining clinical data pair. The value of the latter appears in the bottom cell of the second column of the mini-max table for this diagnosis, and equals the current P of the diagnosis *before* processing clinical data absent. New clinical data with equal or smaller PP value can be disregarded because—even if present—they will not change the current P of this diagnosis; the algorithm moves to Step 4. When the new clinical datum has a PP value that exceeds the current P of the diagnosis before considering clinical data absent (bottom cell of second column), the algorithm recommends it as best cost-benefit clinical datum. The user then verifies whether this clinical datum is absent or present. When this clinical datum is absent, it is disregarded, because if able to change

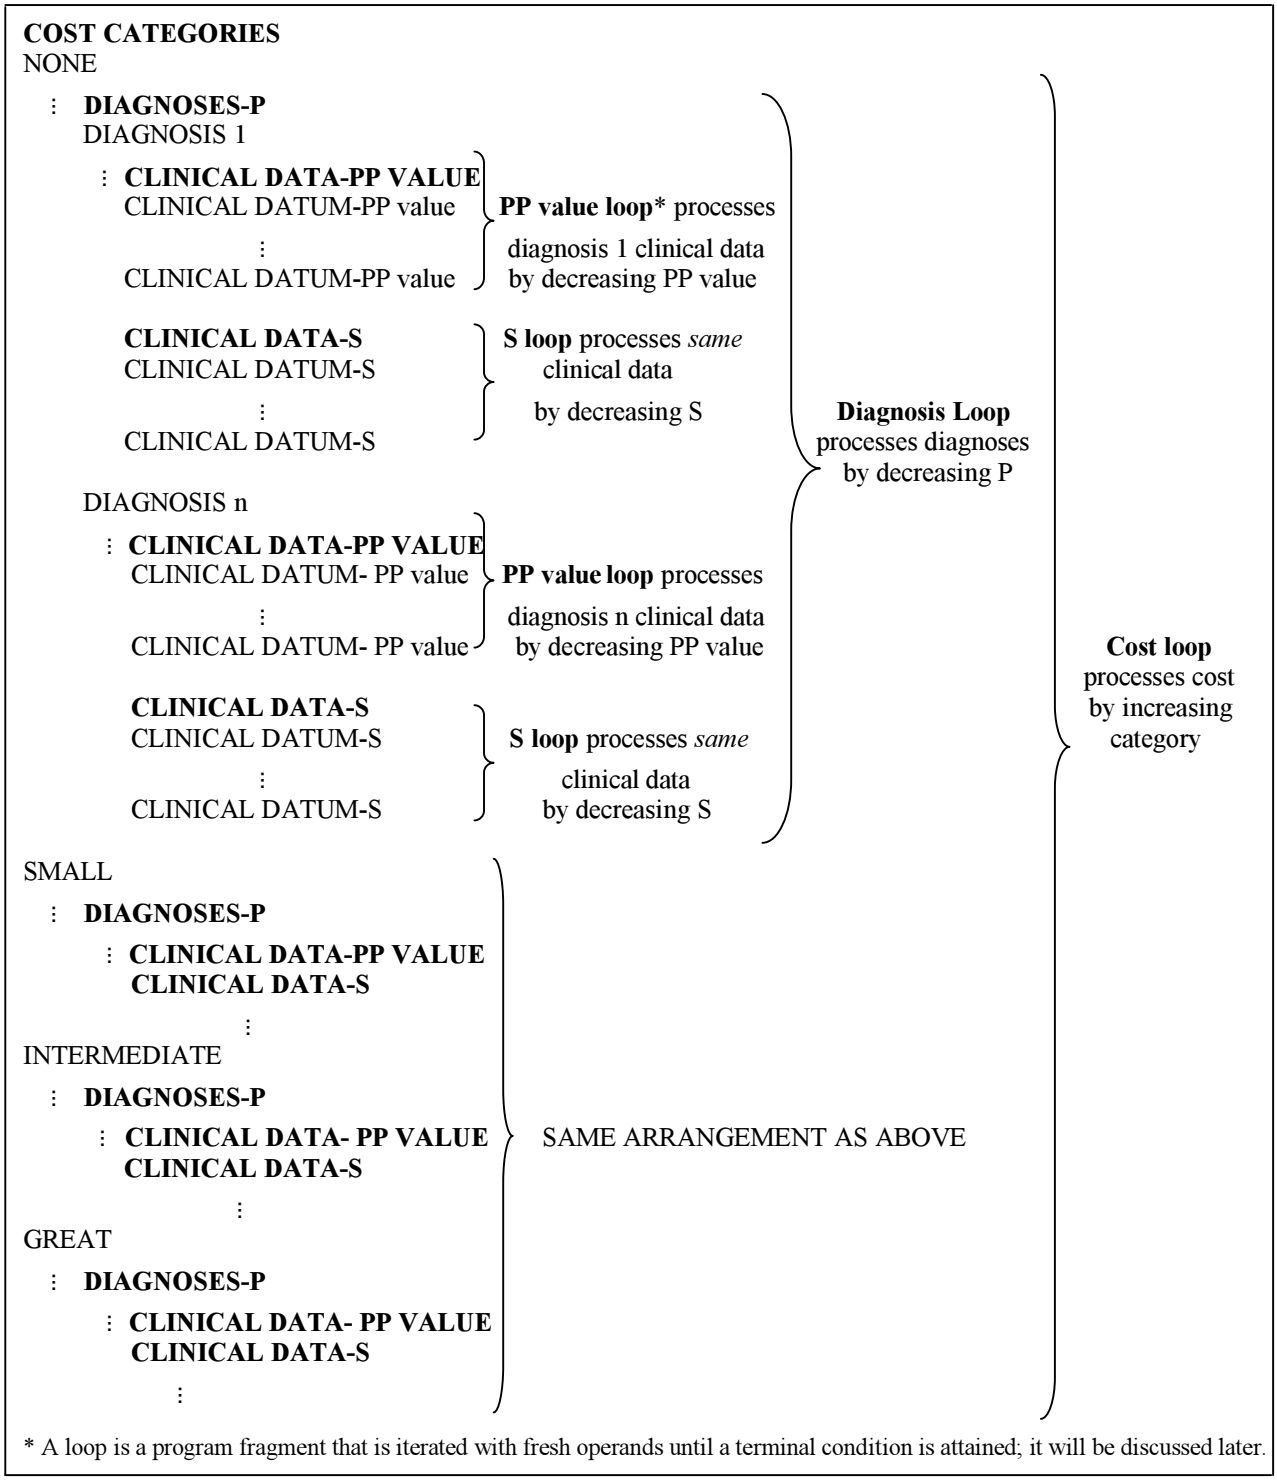

FIGURE 1. Nested loops for selecting best cost-benefit clinical datum.

the total P of the diagnosis, it will be detected by the S loop at the next Step 4, which processes clinical data assumed absent. When the recommended best cost-benefit clinical datum is present, the algorithm generates a new clinical datum list headed by this datum, and the partial P of the diagnosis *before* considering clinical data absent assumes the PP value of this new datum. To recalculate the total P of the

diagnosis *after* considering clinical data absent, several new clinical data pairs, combination of the best cost-benefit clinical datum present with each clinical datum absent, are generated; new clinical data pair tables are created and the partial P values for the diagnosis are calculated. A new row with these values is inserted in each mini-max table and the total P of the corresponding diagnoses are established.

Example: Assume we need to know whether the clinical datum pulmonary mass as evidenced by X-ray plain films, when *present*, can *increase* the current total P of lung cancer (see mini-max table for Lung Cancer.) The PP value of a pulmonary mass for lung cancer, stored in the disease model in the knowledge base, is 0.857. Equation 6 states that the greatest of the clinical data PP values supporting a diagnosis equals the P value of this diagnosis:

$$P_{\text{lung cancer}} = \max (PP\ value_{\text{cough}}, PP\ value_{\text{hemoptysis}}, PP\ value_{\text{dyspnea}}, PP\ value_{\text{expectoration}}, PP\ value_{\text{MTb}}, PP\ value_{\text{mass}})$$

$$= \max (0.241, 0.278, 0.444, 0.104, 0.000, 0.857) = 0.857$$

Accordingly, the previous P of lung cancer (0.444), *before considering clinical data absent* (bottom cell of second column in the lung cancer mini-max table), is increased to its new value of 0.857. The algorithm then creates a new clinical datum list:

Pulmonary mass	**S**	**PP value**
Lung cancer	0.9	0.857
Pulmonary tuberculosis	0.1	0.095
Pulmonary embolism	0.05	0.048
Bronchiectasis	0.0	0.000

To determine the total P of the diagnosis, *after considering clinical data absent*, the algorithm creates two new clinical data pair tables and calculates the partial P values for lung cancer:

Mass-Cavity	**PP value**		**S**			**Partial P**
Lung cancer	0.857	×	(1-0.30) = 0.600	÷	0.686	= **0.875**
Pulmonary tuberculosis	0.095	×	(1-0.60) = 0.038			
Pulmonary embolism	0.048	×	(1-0.00) = 0.048			
Bronchiectasis	0.000	×	(1-0.10) = 0.000			
			0.686			

Mass-Fever	**PP value**		**S**			**Partial P**
Lung cancer	0.857	×	(1-0.10) = 0.771	÷	0.833	= **0.925**
Pulmonary tuberculosis	0.095	×	(1-0.70) = 0.028			
Pulmonary embolism	0.048	×	(1-0.30) = 0.034			
Bronchiectasis	0.000	×	(1-0.00) = 0.000			
			0.833			

In the mini-max table for lung cancer, the algorithm creates a new row that shows these partial P values; then, the total P of this diagnosis *after* considering clinical data absent is calculated:

Mini-max table for lung cancer when pulmonary mass is present

LUNG CANCER	PP value = P before considering absent clinical data	Cavity absent S = 0.3	Fever absent S = 0.1	MINIMUM VALUE IN EACH ROW
Cough present	0.241	0.231	0.297	0.231
Hemoptysis present	0.278	0.254	0.349	0.254
Dyspnea present	0.444	0.402	0.540	0.402
Expectoration present	0.104	0.109	0.135	0.104
MTb present	0.000	0.000	0.000	0.000
Pulmonary mass present	*0.857*	*0.875*	0.925	0.857
MAXIMUM VALUE IN EACH COLUMN	0.857	0.875	0.925	**Total P = 0.857**

MTb, *Mycobacterium tuberculosis*

The total P of **lung cancer** at this diagnostic step is the maximum value (**0.857**) in the last column. Because of pulmonary mass *present* in chest X-ray plain films, the total P of lung cancer increased from 0.402 to 0.857.

Notice that except for the last two rows, partial P values shown in all other cells remain unchanged. However, the value and location of the determining partial P cell changed.

Each new best cost-benefit clinical datum present creates a new clinical datum list that includes the diagnosis from which it was selected. This diagnosis appears in some or all previous clinical datum lists because it originated the search for the new clinical datum; the latter just increases the number of clinical data that support this diagnosis and its P. Some of the new clinical datum lists may include previously unlisted diagnoses that also may manifest this clinical datum (Fig. 2.) When this occurs, such new diagnoses will not have clinical data in common with any previous diagnosis because they were not included in previous clinical datum lists; accordingly, previous and new diagnoses, if confirmed as final, must be concurrent. New clinical datum lists bring up new diagnoses; these, in turn, bring up new clinical data. At first thought, this cycle may seem to iterate indefinitely until the universe of clinical data is exhausted. In reality, this does not occur, because a single patient cannot manifest all clinical data. At some point, the newly recommended best cost-benefit clinical datum will simply be absent and will not create a new clinical datum list, aborting the cycle; still, it must be investigated so as to confirm its absence. Neither can a patient be afflicted by a multitude of concurrent diseases. Another factor limiting the number of diagnoses is that a best cost-benefit clinical datum is selected for its great PP value and accordingly is either pathognomonic for a single diagnosis or supportive of only a few diagnoses.

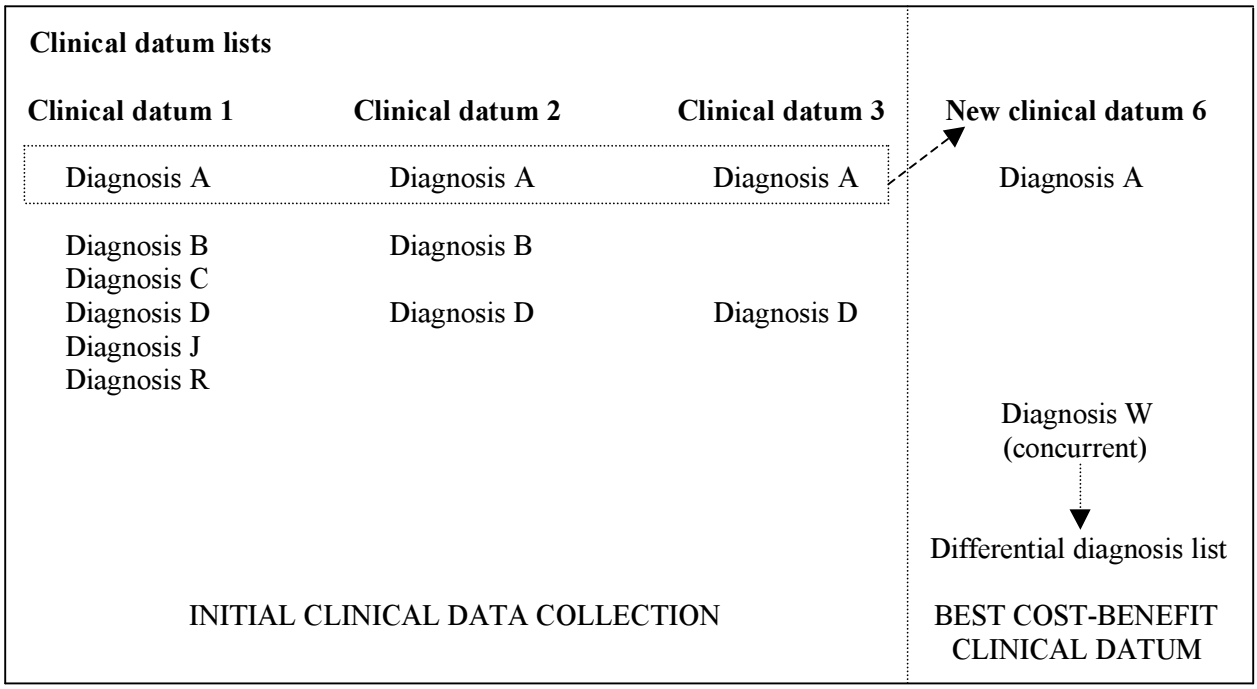

FIGURE 2. To increase probability of diagnosis A, supported by clinical data 1, 2, and 3 (dotted box), clinical datum 6 is recommended as best cost-benefit clinical datum. Clinical datum 6, confirmed present in the patient, creates clinical datum list 6 that includes diagnosis A from which PP list clinical datum 6 was selected; this increases P of diagnosis A and also the number of supporting clinical data. Previously unlisted diagnosis W, also supported by clinical datum 6, is included in the differential diagnosis list as a new diagnosis; if it reaches confirmation threshold, it will become a *concurrent* final diagnosis.

Step 4. Recommend a new clinical datum as best clinical datum assuming it absent

Typically, the greater the S of a clinical datum absent, the more it decreases the total P of a diagnosis. However, with the mini-max procedure, a clinical datum absent can occasionally increase total P (mini-max property 4 B) and it also can occur that a clinical datum with a smaller S can decrease total P more than *another* clinical datum absent with a greater S.

To verify whether a new clinical datum absent will indeed decrease the total P, the mini-max procedure needs not always be applied in its entirety; to save computer time (although this is unlikely to be critical, considering the great speed of today's computers), Tomás Feder devised the 3-Step method (see below). Should the first or second of these steps prove that a clinical datum does not decrease the total P, the following steps would be rendered unnecessary and the algorithm would move to the next diagnosis or cost category.

Besides shortening the processing time, the 3-Step method also determines whether *another* clinical datum absent with smaller S value is nevertheless able to reduce the total P of a diagnosis. This determination is accomplished by the first step only; when a clinical datum absent is thus proved not to reduce the P, neither will another datum of smaller S, and the algorithm can move to the next diagnosis. This is not true for the second and third step; even when these steps prove that the new clinical datum absent does not reduce P, the *next clinical datum with smaller S* must be processed with the first step, to determine whether this smaller S clinical datum can reduce P.

In the same COST category and DIAGNOSIS in which a new best cost-benefit clinical datum assumed present was processed involving PP value list, the algorithm moves now to the S LIST. From this list, the clinical datum of greatest S, not yet selected, is selected and substitutes the clinical datum absent in the *current* determining clinical data pair, creating a new clinical data pair.

3-Step method to determine whether a new clinical datum absent will decrease total P

Step 4.1

In the newly created clinical data pair, equation 7 is applied to the PP value of clinical datum present and S of the new clinical datum from the S list, assumed absent. When the resulting P equals or exceeds the current *total P* of the diagnosis, neither *this* nor any *other* clinical datum in the same S list will decrease the total P; steps 4.2 and 4.3 can be skipped and the 3-Step method can advance to the next diagnosis. Conversely, when the resulting P is smaller, total P of diagnosis under consideration *might* decrease; the resulting P is stored. Proceed with step 4.2.

Computational time: step 4.1 (applying equation 7) involves only one subtraction and one multiplication.

Example: Assume we need to know whether the clinical datum pulmonary mass as evidenced by chest X-ray plain films, when *absent*, can *decrease* the current total P of lung cancer, which is 0.402, as shown in the last cell of the corresponding mini-max table. The P of lung cancer *before* considering clinical data absent was 0.444 (equals greatest PP value in the second column, corresponding to dyspnea present); S of pulmonary mass on X-ray films for lung cancer is 0.9, as shown in the corresponding S list. Substituting these values in equation 7, we obtain:

$$P = PP\ value \times (1 - S) = 0.444\ (1 - 0.9) = 0.044$$

This total P (0.044) is smaller than the prior total P of lung cancer (0.402); therefore, absence of mass in chest X-ray films *might* decrease this P. Proceed with step 4.2.

Step 4.2

A *clinical data pair table* is generated for the newly created clinical data pair of step 4.1; this table is headed by the clinical data pair and lists all the diagnoses in the differential diagnosis list. For each diagnosis, equation 7 is iterated, being applied to the corresponding PP value and S that the heading clinical data pair has for this diagnosis, to obtain the corresponding not yet normalized P (these P values do not sum 1.) Now equation 8 is applied *only* to the diagnosis with the S list from which the clinical datum with greatest S was selected. The numerator of equation 8 is the non-normalized P of this diagnosis and the denominator is the sum of all non-normalized P values corresponding to the diagnoses in the clinical data pair table (this sum is smaller than 1). The value of the numerator of equation 8 was calculated in step 4.1; in the denominator, the clinical data pair for each term is the *same*, but with PP values and S values related to the respective diagnoses. The result is the normalized partial P of the mentioned diagnosis.

When the resulting partial P for the diagnosis under consideration equals or exceeds the current P, neither step 4.2 nor step 4.3 will decrease this total P. When this happens, disregard the clinical datum being processed and iterate the 3-Step method with the next greatest S clinical datum in the S list for the same diagnosis, until step 4.1 proves that no *other* clinical datum with a smaller S can decrease the total P. Then, advance the 3-Step method to the next diagnosis. Conversely, when the resulting partial P is smaller than the current total P, the clinical datum absent *might* decrease total P; the resulting P is stored. Proceed with step 4.3.

Only if the new clinical datum later is confirmed to be absent, will a new column be generated in the mini-max table. At step 4.2, *only one* cell—at the intersection of the clinical datum present and clinical datum absent—of this column (see example below) can show the result of equation 8. This cell shows the partial P that this clinical data pair confers to the diagnosis; all other cells of this column remain blank. Because equation 8 was applied to a clinical data pair with a greater S than the previous one, we know that the resulting partial P is the smallest in its row, but we do not know yet whether it also will be the greatest in its column and become the determining partial P.

Note: step 4.2 involves equation 8, the numerator of which is equivalent to equation 7 and the denominator of which always is equal or smaller than 1. Accordingly, the result of equation 8 (partial P of the diagnosis in step 4.2) will always be greater than its numerator (P of the diagnosis in step 4.1.) For this reason, it is unnecessary to proceed with step 4.2 when step 4.1 already yielded a P that equals or exceeds prior P.

Computational time: step 4.2 involves the application of equation 7 to *all diagnoses* in *only one clinical data pair table*, and equation 8 to *only one diagnosis* in this table. Computational time equals step 4.1 (equation 7) times the number of diagnoses (denominator of equation 8) plus the sum of the resulting terms (one sum for each diagnosis) plus the division of the numerator by the denominator.

Example (continued): create the dyspnea-mass clinical data pair table with the differential diagnoses showing the respective PP values and S values and then apply equation 7 to each diagnosis:

Dyspnea-Mass	**PP value**	**S**
Pulmonary tuberculosis	0.148 ×	(1- 0.10) = 0.133
Pulmonary embolism	0.370 ×	(1- 0.05) = 0.352
Bronchiectasis	0.037 ×	(1- 0.00) = 0.037
Lung cancer	0.444 ×	(1- 0.90) = 0.044

Then apply equation 8 *only* to the diagnosis being processed, to obtain the partial P that this clinical data pair confers to this diagnosis:

$$P_{\text{lung cancer}} = \frac{\text{PP value}_{\text{lung cancer}}\,(1-S_{\text{lung cancer}})}{\text{PP value}_{\text{lung cancer}}\,(1-S_{\text{lung cancer}})+\text{PP value}_{TB}\,(1-S_{TB})+\text{PP value}_{\text{embolism}}\,(1-S_{\text{embolism}})+\text{PP value}_{\text{bronchiectasis}}\,(1-S_{\text{bronchiectasis}})}$$

$$P_{\text{lung cancer}} = \frac{0.444\,(1-0.90)}{0.444\,(1-0.90)+0.148\,(1-0.10)+0.370\,(1-0.05)+0.037\,(1-0.00)}$$

$$= \frac{0.044}{0.044+0.133+0.352+0.037} = 0.078$$

Dyspnea-Mass	**PP value**	**S**		**Partial P**
Pulmonary tuberculosis	0.148 ×	(1- 0.10) = 0.133		
Pulmonary embolism	0.370 ×	(1- 0.05) = 0.352		
Bronchiectasis	0.037 ×	(1- 0.00) = 0.037		
Lung cancer	0.444 ×	(1- 0.90) = 0.044	÷ 0.566 = 0.078	

We now can include in the mini-max table the new column "Mass absent", listing only one partial P conferred by the dyspnea-mass clinical data pair to lung cancer diagnosis; the other cells of the column remain empty.

Mini-max table for lung cancer when pulmonary mass is absent

LUNG CANCER	PP value = P before considering absent clinical data	Cavity absent S = 0.3	Fever absent S = 0.1	Mass absent S = 0.9	MINIMUM VALUE IN EACH ROW
Cough present	0.241	0.231	0.297		
Hemoptysis present	0.278	0.254	0.349		
Dyspnea present	0.444	0.402	0.540	0.078	0.078
Expectoration present	0.104	0.109	0.135		
MTb present	0.000	0.000	0.000		
MAXIMUM VALUE IN EACH COLUMN	0.444	0.402	0.540		

MTb, *Mycobacterium tuberculosis*

Because 0.078 is smaller than 0.402 (the current total P), the possibility remains that this total P *might* be decreased. Proceed with step 4.3.

Step 4.3

All possible new clinical data pairs and their tables are created, with each clinical datum present combined with the new clinical datum absent. In step 4.2, equation 7 was applied to all diagnoses in only one clinical data pair table, and then equation 8 was applied only to the diagnosis being processed, but now in step 4.3, equation 8 is applied to the same diagnosis in *all* the clinical data pair tables. This calculates the partial P values that the new clinical data pairs confer to this diagnosis; these partial P values show in *all* cells of the new column of the mini-max table. This enables to determine the total P of the diagnosis, showing whether the absent clinical datum of greatest S indeed decreases it. If so, this clinical datum is recommended as the best cost-benefit clinical datum next to investigate; otherwise it is disregarded and the algorithm moves to the next diagnosis.

Computational time: equals the duration of step 4.2 multiplied by the number of new clinical data pairs.

Example (continued): create the remaining new clinical data pair tables and iterate equation 8 for lung cancer in all of them:

Dyspnea-Mass	**PP value**	**S**			**Partial P**
Pulmonary tuberculosis	0.148	× (1- 0.10)	= 0.133		
Pulmonary embolism	0.370	× (1- 0.05)	= 0.352		
Bronchiectasis	0.037	× (1- 0.00)	= 0.037		
Lung cancer	0.444	× (1- 0.90)	= 0.044	÷ 0.566	= 0.078
			0.566		
Cough-Mass					
Pulmonary tuberculosis	0.276	× (1-0.10)	= 0.248		
Pulmonary embolism	0.172	× (1-0.05)	= 0.164		
Bronchiectasis	0.310	× (1-0.00)	= 0.310		
Lung cancer	0.241	× (1-0.90)	= 0.024	÷ 0.747	= 0.032
			0.747		
Hemoptysis-Mass					
Pulmonary tuberculosis	0.222	× (1-0.10)	= 0.200		
Pulmonary embolism	0.333	× (1-0.05)	= 0.317		
Bronchiectasis	0.167	× (1-0.00)	= 0.167		
Lung cancer	0.278	× (1-0.90)	= 0.028	÷ 0.711	= 0.039
			0.711		

Expectoration-Mass

Pulmonary tuberculosis	0.417	× (1-0.10)	= 0.375			
Pulmonary embolism	0.010	× (1-0.05)	= 0.010			
Bronchiectasis	0.469	× (1-0.00)	= 0.469			
Lung cancer	0.104	× (1-0.90)	= 0.010 / 0.864	÷ 0.864	= 0.012	

Mycobacterium TB-Mass

Pulmonary tuberculosis	1.000	× (1-0.10)	= 0.900			
Pulmonary embolism	0.000	× (1-0.05)	= 0.000			
Bronchiectasis	0.000	× (1-0.00)	= 0.000			
Lung cancer	0.000	× (1-0.90)	= 0.000 / 0.900	÷ 0.900	= 0.000	

We now can complete the new column "Mass absent" in the mini-max table for lung cancer:

Mini-max table for lung cancer when pulmonary mass is absent

LUNG CANCER	PP value = P before considering absent clinical data	Cavity absent S = 0.3	Fever absent S = 0.1	Mass absent S = 0.9	MINIMUM VALUE IN EACH ROW
Cough present	0.241	0.231	0.297	0.032	0.032
Hemoptysis present	0.278	0.254	0.349	0.039	0.039
Dyspnea present	0.444	0.402	0.540	*0.078*	0.078
Expectoration present	0.104	0.109	0.135	0.012	0.012
MTb present	0.000	0.000	0.000	0.000	0.000
MAXIMUM VALUE IN EACH COLUMN	0.444	0.402	0.540	0.078	**Total P = 0.078**

MTb, *Mycobacterium tuberculosis*

Were a pulmonary mass *absent*, the previous total P of lung cancer indeed would be reduced from 0.402 to **0.078**.

A clinical datum absent with a smaller S occasionally is able to reduce total P more than another with a greater S. This mandates processing the clinical datum in the S list immediately below that with greatest S, because this clinical datum, if absent, might further reduce P. Accordingly, the 3-Step method is iterated with clinical data with progressively smaller S until step 4.1 no longer reduces P. Manually processed examples demonstrated that such reduction of total P occurs only occasionally and ceases after one or a few clinical data with smaller S are processed, and does not significantly prolong selection of the best cost-benefit clinical datum. This is because only clinical data with slightly smaller S can additionally reduce the total P.

***

So far, we explained how a best cost-benefit clinical datum assumed *absent* is selected and recommended. Now, this best cost-benefit clinical datum must be investigated for actual presence or absence in the patient. When this clinical datum is present, it is disregarded, because it neither can decrease the total P of the diagnosis (general rule) nor increase the total P (were it able to do so, it would have been detected by the PP value loop in step 3.) When this clinical datum is absent, it possibly could be masked by a drug or concurrent disease interaction. If this clinical datum absent is flagged with an interaction identifier, the algorithm must check for a drug or disease able to mask it; if confirmed, the clinical datum absent is disregarded and the user notified. If the clinical datum absent is not masked, the

greatest partial P in the new column in the mini-max table becomes the new determining partial P, and the total P of the diagnosis decreases to this value.

In summary, to select the best cost-benefit clinical datum to investigate next, the algorithm loops at three nested levels (Figure 1): outer, intermediate, and inner. (1) The outer *cost loop* processes clinical data not yet investigated in order of increasing cost category: none, small, intermediate, and great. (2) Within each cost category, the intermediate *diagnosis loop* processes the diagnoses of the differential diagnosis list in order of decreasing P because those with greatest P values, are the best candidates for a final diagnosis, and can sooner conclude the diagnostic quest. (3) The inner *clinical data loop* comprises two sub-loops: the first begins at the top of the PP value list and terminates when no clinical datum exists with a PP value greater than the PP value of the determining clinical data pair. The second sub-loop processes, within the same level, the clinical data sorted by decreasing S, terminating when no clinical datum exists, able to change the P of the corresponding diagnosis.

All diagnoses in each cost category are similarly processed. The user is prompted each time the loop goes to a greater cost category. The remaining differential diagnoses with their P are displayed and the user is asked whether he wants to proceed in the greater cost category or prefers a deferred diagnosis, diagnosis by exclusion, or empirical treatment. The entire looping process terminates when all final diagnoses are obtained, the cost exceeds the benefit, or all the clinical data able to change P of diagnoses are processed. Clinical data that have the greatest PP value or the greatest S typically involve costly pathological investigations, such as biopsy or even autopsy. To request a biopsy or even an autopsy for a patient with tonsillitis would make little sense. This exaggeration emphasizes the importance of initially considering the cost of a clinical datum, before evaluating its PP value or S. However, in an emergency or when a patient's condition is deteriorating, investigation of confirmatory clinical data of great PP value takes priority over cost.

The recommended best cost-benefit clinical datum could be a common symptom quickly asked or immediately observed by the physician. Should obtaining a clinical datum require an involved test or procedure, the diagnostic process must be interrupted until the result becomes available. The "position of the game board", so to say, must be saved in the computer and opportunely retrieved to continue the "game", because each new clinical datum, with its presence or absence in the patient, sets a new stage for the selection of the next best cost-benefit clinical datum. A disease is not a static process. If the clinical picture changes considerably prior to obtaining the diagnostic procedure result, a new diagnostic evaluation must be accomplished, sometimes from the beginning.

VII. SIMULTANEOUS RECOMMENDATION OF BEST COST-BENEFIT CLINICAL DATA

Few diagnostic computer programs recommend a single best cost-benefit clinical datum next to investigate; I know of none that simultaneously recommends a **set** of such data.

Recommending one best cost-benefit clinical data at a time is impractical; were the patient to require several tests, he would need to be contacted repeatedly for additional instructions. For this reason, physicians often *simultaneously* order a set of several analyses, tests, or procedures; this approach is critical for emergency cases, where there is no time to wait for sequential results.

A computer can emulate such human behavior by iterating the best cost-benefit clinical data function, first assuming each newly recommended best cost-benefit clinical datum as virtually present and then as virtually absent, while observing the effect that each iteration has on the P of each diagnosis. Such iteration can be represented as a trichotomy tree (Fig. 3.)

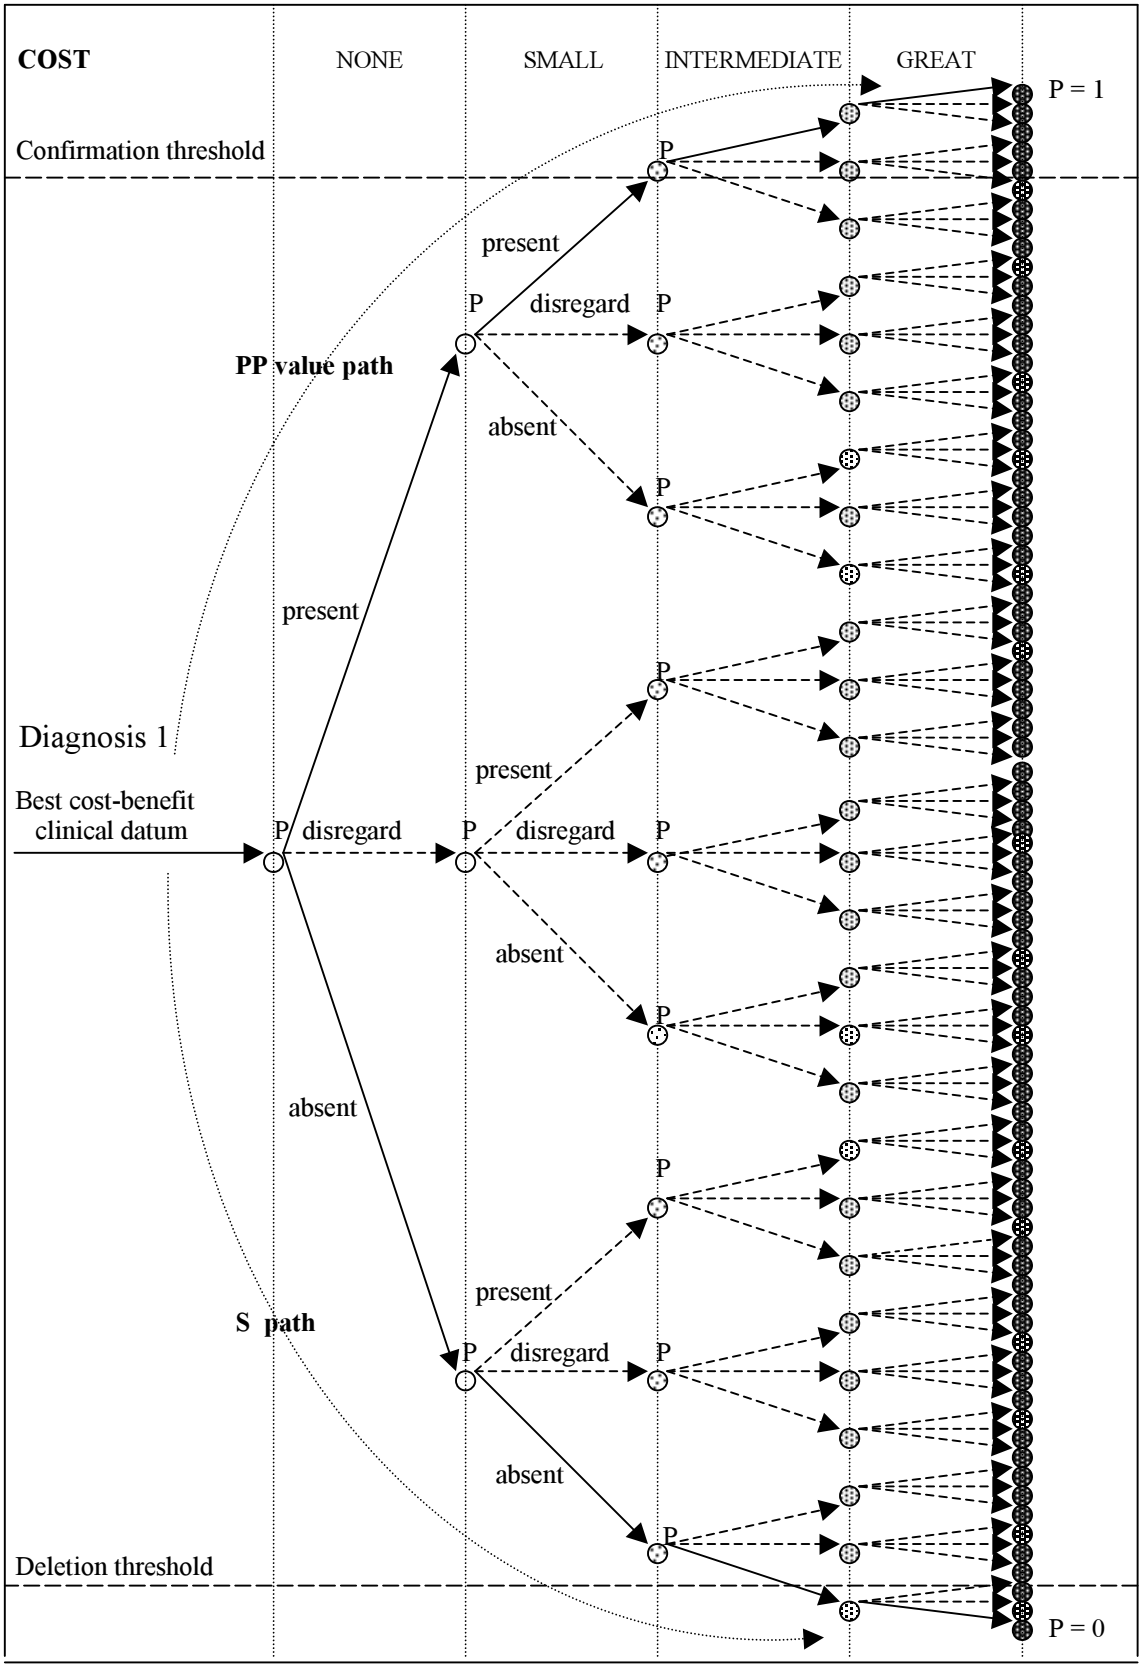

FIGURE 3. Trichotomy tree. Each branch (arrow) represents a best cost-benefit clinical datum. Each node (circle) represents the cumulative P and cost of the diagnosis.
○ NO COST; ◌ SMALL COST; ◍ INTERMEDIATE COST; ● GREAT COST

Each tree represents a single diagnosis in the differential diagnosis list. Virtual branches represent best cost-benefit clinical data present or absent; nodes represent the P and cost of the diagnosis. Each node originates three new branches; four cost level iterations are involved; accordingly, the total number of branches is $3^1 + 3^2 + 3^3 + 3^4 = 120$. Each top branch originating at a node assumes that the best cost-benefit clinical datum—selected from the PP value list—is present; accordingly, it increases P of the diagnosis and is depicted by an ascending arrow. Each middle branch assumes that *this* best cost-benefit clinical datum is absent; accordingly, it is disregarded, does not change P, and is depicted by a horizontal arrow. Each bottom branch assumes that the best cost-benefit clinical datum—selected from the S list—is absent; accordingly, it decreases P and is represented by a descending arrow. The same middle branch also assumes that this best cost-benefit clinical datum, selected from the S list, is present, disregarded, and does not change P. Additionally, this middle branch represents situations in which no best cost-benefit clinical datum was found in either the PP value list or S list; accordingly, it offers no best cost-benefit clinical data to investigate. This reduces the best cost-benefit clinical data to investigate to only two branches per node, the top (present) and the bottom (absent) branches; the middle branch is preserved however, because it takes us to a next node. In the entire tree, the number of best cost-benefit clinical data to investigate (branches) now is reduced to 80. This result multiplied by the number of diagnoses (trees) in the differential diagnosis list yields the total number of best cost-benefit clinical data to investigate, provided no best cost-benefit clinical data are shared with other trees (diagnoses.)

A clinical datum frequently is shared by diverse diagnoses. However, a best cost-benefit clinical datum selected from the PP value list has a great PP value; and when present, it is either very characteristic for only a few diagnoses or pathognomonic for only one diagnosis. Such clinical data is typically not shared with any other tree. When a best cost-benefit clinical datum is selected from the S list, because of its great S, and is absent, it might be shared by diverse diagnoses for which this clinical datum occurs frequently.

In a single tree, the best cost-benefit clinical datum present, represented by the top branch, typically has a great PP value and strongly supports a diagnosis. The best cost-benefit clinical datum absent, represented by the bottom branch, typically has a great S and strongly opposes a diagnosis. Occasionally, a best cost-benefit clinical datum, when it has a great PP value *and* a great S, may be recommended simultaneously in the top and bottom branch. This apparent opposition is not conflicting because it refers to *virtually present and absent alternatives* that do not coexist in a real patient case.

Processing the entire **set** of best cost-benefit clinical data perhaps could at once confirm as final those diagnoses with a P close to 1 and rule out those with a P close to 0. However, exhaustively traversing all branches of this exponentially growing trees is limited by the increasing number of clinical data to investigate and the cost involved. At best, the trichotomy tree approach may enable us to move only a few steps forward. Fortunately, heuristic shortcuts might dispel this concern. As noted earlier, clinical data present of great PP value that strongly favor a diagnosis are unlikely to be opposed by clinical data absent of great S that strongly disfavor the same diagnosis. Accordingly, when a diagnosis with an initial great P is processed, we would expect the algorithm to recommend a best cost-benefit clinical datum of greater PP value that would further enhance that P, rather than recommend a best cost-benefit clinical datum of great S. Conversely, when a diagnosis with an initial small P is processed, we would expect the algorithm to recommend a best cost-benefit clinical datum of greater S that would further reduce that P, rather than recommend a best cost-benefit clinical datum of great PP value. This expectation would favor virtual traversing from present to present branches toward greater P values of the diagnosis or from absent to absent branches toward smaller P values of the diagnosis. Thus, in the tree of Fig. 3 the process would tend to traverse solid exterior branches only, while avoiding zigzag traversal along dashed alternating present and absent interior branches.

If we elect not to *exclusively* traverse extreme exterior branches, a few virtual best cost-benefit clinical data alternatively absent and present can be accepted. This maintains traversal *near* the exterior branches, leading to nodes with P near to diagnosis confirmation or elimination values.

Ideally, the algorithm explores—for each diagnosis in the differential diagnosis list—all possible virtual traversals until maximum or minimum P are attained, or until all available clinical data are exhausted. To accomplish this goal, prompts and authorization requests to continue in the next greater cost category must be bypassed. Cost momentarily is disregarded so as to obtain an ample overview of all best cost-benefit clinical data available. A decision regarding which best cost-benefit clinical data to select and up to what cost can be made afterwards according to disease severity and other circumstances.

NO COST clinical data are obtained as part of the initial consultation, from the history and physical examination. After evaluation of these NO COST clinical data, the user follows one or more of diverse strategies to select the most appropriate set of best cost-benefit clinical data next to investigate:

- Select *all* clinical data shown in the tree; this is possible only if the cost of acquiring them does not surpass an acceptable limit.
- Select only those clinical data of *lesser cost*. If, as a consequence, the diagnostic quest is not concluded, investigation of costlier clinical data can be pursued at a subsequent consultation. This strategy is consistent with standard medical practice.
- Select only best cost-benefit clinical data corresponding to *exterior branches* (represented by the solid arrows in Fig. 3) which traverse to a P that equals or approximates 1 or 0.
- Select preferably best cost-benefit clinical data *shared* by more than one diagnosis, when available; such data can be used for calculating P of more than one diagnosis.

Such selections can be made by the algorithm if parameters of cost, number of tests, and computational time limits can be provided.

The validity of the tree must be confirmed by comparing a prototype computer program with actual patient cases.

VIII. CONCLUSION OF DIAGNOSTIC QUEST

How does the human mind or the computer detect that the diagnostic process is completed and that a final diagnosis best represents the disease that afflict the patient? Several methods can determine when the conclusion of the diagnostic quest is achieved; our algorithm uses a *double threshold method*.

Double threshold method establishes two empiric thresholds: a *confirmation threshold* and a *deletion threshold*. When the P of a diagnosis in the differential diagnosis list is greater than the confirmation threshold, such diagnosis is flagged as *final*. When the P of a diagnosis in the differential diagnosis list is smaller than the deletion threshold, such diagnosis is flagged as *deleted* (ruled out.) The diagnostic quest concludes when all diagnoses in the differential diagnosis list are flagged. Flagging and maintaining final and deleted diagnoses in the differential diagnosis list, as opposed to removing them, keeps the number of diagnoses equal to the number of terms in the denominator of equation 8 and equal to the number of mini-max tables, avoiding recalculations of partial P values (a detailed discussion of this matter can be found in our book [6].) Final diagnoses are copied into a *final diagnoses list* and displayed together

with the clinical data that support them. The double threshold method, based on the mini-max procedure, enables the identification and separation of concurrent from competing diagnoses.

Each concurrent final diagnosis requires a "clinching" clinical datum present (with PP value close to 1); very sensitive (S close to 1) clinical data absent rule out diagnoses. For final diagnoses, P = 1; for deleted diagnoses, P = 0. In actual practice, such extreme values are not required; the confirmation and deletion threshold values must be empirically determined.

Threshold values should be adjustable according to the severity of patient's condition, favorable or unfavorable evolution, prognosis, and availability of efficient treatments. This problem may be related to a type of *Receiver Operating Characteristic (ROC) curve* that can be used to tune the diagnostic process. Decreasing a confirmation threshold may erroneously confirm a diagnosis of a disease that does not afflict the patient; increasing it precludes this error, but may require costlier clinical data to increase the P of the diagnosis to the level of this threshold. Increasing a deletion threshold may erroneously rule out a diagnosis of a disease that afflicts the patient; decreasing it precludes this error, but may require costlier clinical data to decrease the P of the diagnosis, to the level of this threshold.

IX. COMPLEX CLINICAL PRESENTATIONS AND THEIR MODELS

Throughout the interaction between the physician and patient, involving medical examination, tests, and procedures, a jumble of clinical data is produced, which must be coalesced and organized into final diagnoses.

The diagnostic process comprises several levels of complexity. Related clinical data cluster to a syndrome, simple syndromes comprising only a few clinical data coalesce to a complex syndrome or disease, and sometimes to a yet more complex clinical presentation, where the relation of clinical data becomes less obvious.

The algorithm thus far presented uses **probabilistic** calculations, with mini-max procedure, best cost-benefit clinical datum, and discrimination between competing diagnoses and concurrent diagnoses, to determine the P of a final diagnosis. It will work well with simple clinical entities, such as uncomplicated diseases or syndromes where clinical data typically are interrelated and linked to a single cause or lesion. Examples of such simple diseases or clinical entities include bronchitis, asthma, gastroenteritis, hyperthyroidism, obstructive jaundice, and renal failure. At this diagnostic stage, a single final diagnosis accounts for all manifested clinical data. The accuracy of P depends on the accuracy of the S of clinical data. Determination of S for all clinical data and diseases, with equation 2, can be accomplished by (a) retrospectively incorporating information from past medical records and publications; (b) prospectively incorporating information from future medical records. The latter task would be an ongoing project, but would be facilitated by universal electronic storage and processing of medical records; (c) basing the sensitivities on the personal estimation of clinicians and specialists, but this is subjective and inaccurate. In any case, the task will be labor intensive and require the cooperation of a medical team. However, probabilistic calculations are feasible as long as they are applied to simple diseases or clinical entities, and a statistically significant number of cases is available. The computational time of this stage is expected not to be excessive (non-exponential.)

In an actual patient, the clinical picture might be more complicated; as a fact, severely ill patients in intensive care units often have multi-organ involvement, present multiple and proteiform clinical data, and may mandate consultation with several specialists. For example, coronary artery disease, acute myocardial infarction, congestive heart failure, shock, *and* thromboembolism in a single patient. A specific disease can manifest diverse clinical forms and clinical presentations, complicating the diagnostic

process. This situation makes impossible to determine the S of each clinical datum for the entire *complex clinical presentation* that involves multiple clinical forms, concurrent diseases, and multiple pathogenic and pathophysiologic mechanisms. It would require analyzing a statistically significant number of cases with identical combinations of clinical entities; it also would take us into an exponential or NP-complete computational time and complexity. Accordingly, probabilistic methods are unsuitable for processing *any* complex clinical presentation; indeed, to my knowledge, no commercial diagnostic programs that can accomplish this exist. A **categorical** method for processing complex clinical presentations is mandatory.

For this reason our algorithm, with its heuristic principles and moderate use of probability, diagnoses first only relatively simple syndromes, clinical entities, and diseases. Let the diagnostic algorithm produce as many final diagnoses of simple concurrent diseases, syndromes, complications, etc. as the clinical data dictate. Our algorithm is able to diagnose satisfactorily these simple clinical entities and also to recognize concurrency. Once we have these partial components (clinical entities), the knowledge base must offer categorical models (*clinical presentation models*; see below), one for each possible clinical interrelation or association of these entities: causal relations, stages, complications, severity, type, localization, etc. We believe that such clinical presentation models, although numerous, are not excessive, and are described in any authoritative medical textbook. The algorithm selects the clinical presentation model that accommodates best all the final clinical entities and diagnoses, probabilistically determined. It reminds us the jigsaw puzzles where different pieces must fit together in some meaningful way. This diagnostic stage does not require a probabilistic approach, but a pure categorical one.

Categorically relating clinical entities based on their associated pathophysiologic links or statistical correlations, into a complex clinical presentation mandate creating a specific model for each possible combination. We call these models complex clinical presentation models.

Complex clinical presentation model: a list of all diseases or clinical entities that are potentially related by pathophysiologic links or statistical correlations (examples in Fig. 4, below.)

These examples are incomplete; not all linked clinical entities are shown. The direction of a causal link, when known, is indicated by an arrow. A complex clinical presentation model should comprise only clinical entities that present a close pathophysiologic relationship or clear statistical correlation. When a clinical entity (*e.g.*, atrial fibrillation in examples 1 and 2) appears in diverse complex clinical presentation models (*e.g.*, atrial fibrillation linked to myocardial infarction in one complex presentation model and linked to hyperthyroidism in another), such clinical entities should *not* be cross-linked. Dissimilar pathophysiologic mechanisms for such clinical entities may be involved in diverse presentation models. Such "trans-model" links (dotted lines in Fig. 4) would result in undesirable entangled networks, as encountered in some diagnostic programs, and which greatly complicate implementation, utilization, and updating. Of all clinical entities that appear in the complex presentation models, only final diagnoses (italicized in Example 1) determined in an actual patient are displayed at the conclusion of the diagnostic quest.

A *complex clinical presentation model* comprises related *clinical entities* and *diseases*; clinical data are excluded from this definition because they are elements of a *disease model*.

Integration of each complex clinical presentation model requires searching of the mentioned relationships in medical literature.

Examples of complex clinical presentation models:

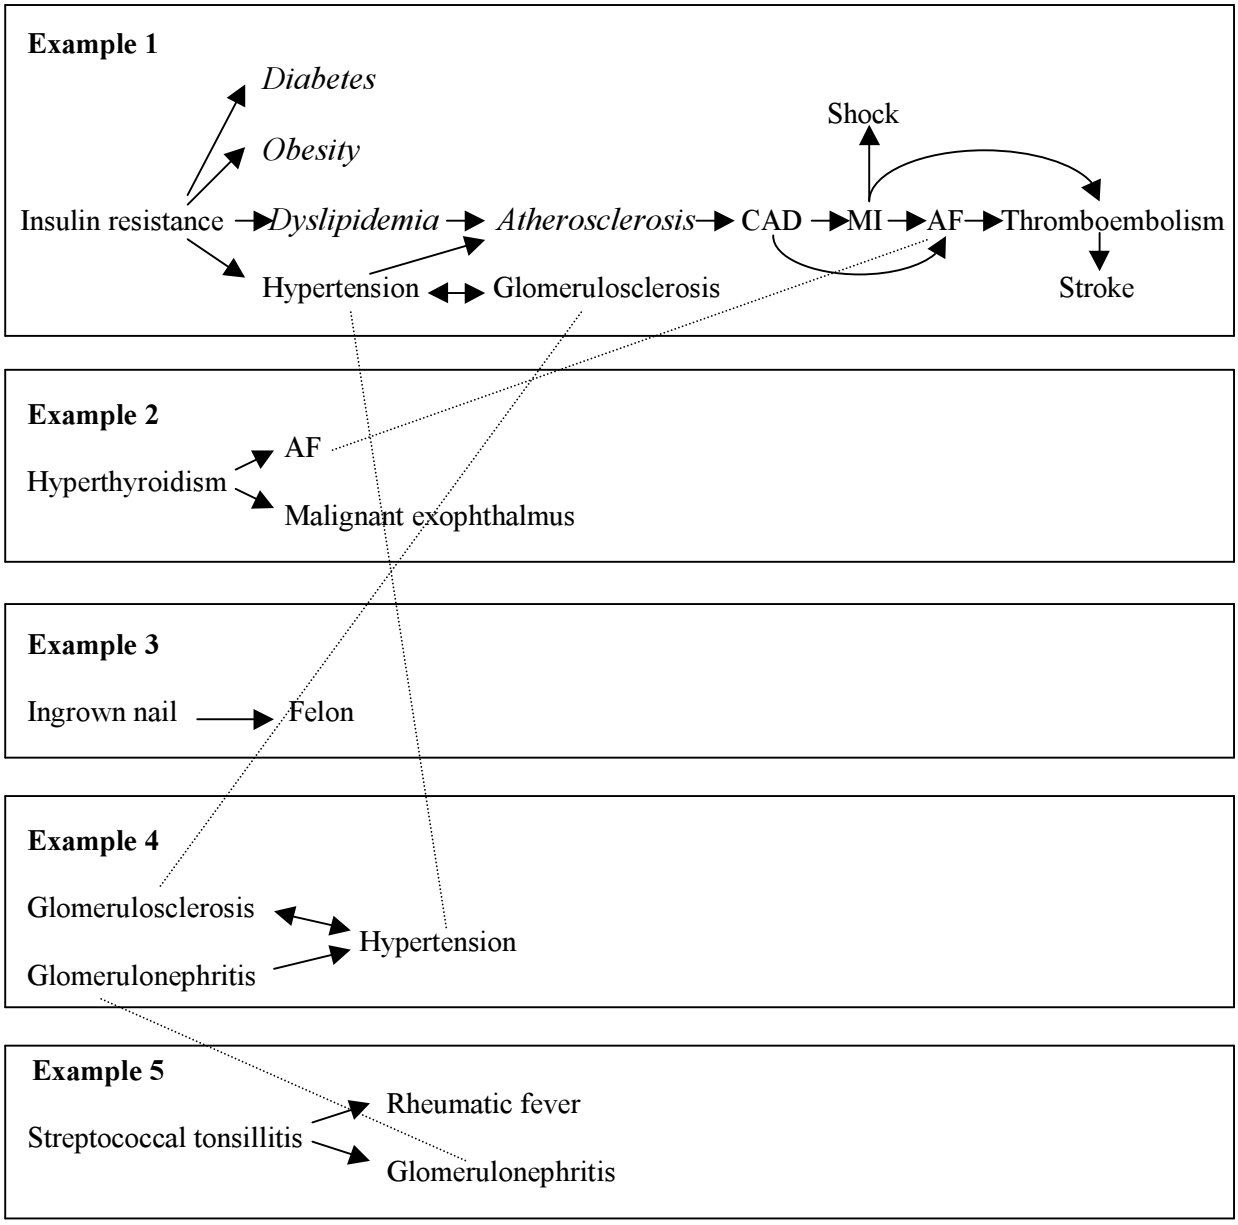

FIGURE 4. CAD, coronary artery disease; MI, myocardial infarction; AF, atrial fibrillation. Dotted lines represent *forbidden* cross-links between complex clinical presentation models (see text).

Should two or more final diagnoses be obtained, the question arises as to whether the respective diseases are related or unrelated. When a relationship cannot be established the diseases are *unrelated concurrent*. Example: tonsillitis and a uterine fibroid are unrelated concurrent diseases, because no known pathogenesis or pathophysiology relates them. When either a statistical correlation (mechanism unknown) or pathophysiologic relationship can be established, we call these diseases *related concurrent* because actually they represent a single complex disease. An example is the metabolic syndrome that associates diabetes, obesity, dyslipidemia, hypertension, and vascular lesions. A statistical correlation among these diseases initially was observed;

subsequent discovery of an underlying pathophysiologic common denominator (insulin receptor resistance) unified the apparently unrelated diseases to a complex clinical presentation. The knowledge base must be provided with *complex clinical presentation models* ideally listing all published *categorical* (as opposed to probabilistic) relations among clinical entities. When the final diagnoses of clinical entities fit into some of such complex clinical presentation models, they represent related concurrence; otherwise, they represent unrelated concurrence.

Several algorithms employ tree and network structures that extend from cause of disease to clinical data and *vice versa*, through several layers or levels, placing probabilities on nodes, branches, and leaves. Most such structures are complex and required years to assemble. I suspect that such structures are difficult to update and would need to be redesigned every few years. In contrast, our algorithm can relatively easily be updated at any time, by simply updating in disease models the S values of clinical data, adding or deleting clinical data when necessary, or adding or deleting disease models. Occasionally, the list of clinical entity combinations in the complex clinical presentation models must be modified when new associations are discovered. However, because the chain of clinical events linking cause of disease, lesions, clinical data, syndromes, and clinical presentation exists, we cannot completely ignore the tree or network structure. When required, as in clinical presentation models, we use these structures categorically, avoiding statistical and probabilistic calculations along their nodes, branches, and leaves. Complex clinical presentation models can be conceived as lists of causally or hierarchically related clinical entities, and not necessarily as trees or networks; in either form, we insist that they must be processed categorically, without any probabilistic component. A relational database can convert ("normalize") a tree or network to a list and *vice versa* [8].

In summary, the entire diagnostic process is achieved in two steps:

Step 1. Probabilistic processing of clinical data

This step matches patient clinical data with disease model clinical data, yielding a differential diagnosis list. Mini-max procedure, best cost-benefit clinical data next to investigate, safety checks for interaction and risk identifiers, and discrimination between competing diagnoses and concurrent diagnoses achieve as many concurrent final diagnoses of clinical entities as are required to account for all manifested clinical data. Then, the algorithm proceeds to step 2.

Step 2. Categorical processing of clinical entities

Step 2.1. Safety checks to avoid overlooking related clinical entities: The algorithm applies a safety check that compares each final diagnosis (*e.g.*, pneumothorax) to all complex clinical presentation models in the knowledge base. Clinical entities related to pneumothorax (*e.g.*, tuberculosis, emphysema, or lung cancer) occurring in matching models are included in the differential diagnosis list to be processed for their presence or absence in the patient.

Step 2.2. Establishing related or unrelated concurrence of clinical entities: When two or more clinical entities are confirmed as final diagnoses, they are compared to all clinical presentation models in the knowledge base. When no complex clinical presentation model linking two or more clinical entities exists, such clinical entities are displayed as unrelated concurrent diagnoses, and the diagnostic quest concludes. When such a complex clinical presentation model exists, the linked clinical entities are displayed as related concurrent, and the diagnostic quest concludes.

Several situations may occur when matching and selecting complex clinical presentation models:

- When more than one clinical presentation model is matched, that which links the greatest number of confirmed final diagnoses will be selected to account for the complex clinical presentation.
- When two or more clinical presentation models link a like number of confirmed final diagnoses, all such models will be selected as alternative representations of the complex clinical presentation.
- When diagnosed clinical entities match only some entities of the complex clinical presentation model, this nevertheless suffices to select the model and link the clinical entities. For example, not all complications listed in a single complex clinical presentation model need be matched; one complication (*e.g.*, pneumothorax) and one cause (*e.g.*, emphysema) suffice to link these two clinical entities.
- When one set of diagnosed clinical entities matches one clinical presentation model, while another set matches another clinical presentation model, the sets do not share clinical entities and the two clinical presentations are concurrent.
- When clinical presentation models involve clinical entities that are related in opposite directions (*e.g.*, myocardial infarction → shock; or shock → myocardial infraction), such alternatives are displayed and the user selects the direction that is most consistent with the clinical picture.

COMMENTS

Beyond the discussion of mini-max procedure, our complete diagnostic algorithm includes several other important functions that are described in our book *Computerized Medical Diagnosis: A Novel Solution to an Old Problem* [6.] Some of these functions, not discussed in the present paper are: interaction identifiers to detect potential disease and drug interactions, risk identifier to minimize the possibility of missing a diagnosis because of failure to collect adequate clinical data or to include the correct diagnosis in the differential diagnosis list, diagnoses related to patient sex and other attributes, diagnosis of somatization and malingering, empirical treatment, deferred diagnosis, and diagnosis by exclusion. The mentioned book also discusses in greater detail the mini-max procedure and presents a simplified version of this procedure.

In recent years, medical applications of computer research have been focused on algorithms dealing with time-related clinical data [9]. To establish a possible relationship (causal or stages of disease) between two clinical events, it is necessary to know the order of occurrence and the elapsed time interval. Because our algorithm operates in a cross-sectional manner, freezing time at a specific diagnostic moment, we have the impression that time is not critical, although significant time or sequence relations of clinical data or diagnoses can be integrated in the disease models and clinical presentation models, where appropriate.

NOVEL CONCEPTS OF OUR ALGORITHM

Our diagnostic program includes novel ideas not mentioned in the extensive literature references reviewed. Some of the ideas are original in themselves; others are original based on a special manner in which extant elements are combined.

- *Disregarding disease prevalence* is doubly advantageous: (1) Prior probabilities of diseases (equivalent to prevalence) are eliminated from Bayes formula, which is transformed into a simplified equation 5 for calculating the PP value from statistically established S values of clinical data. (2) Low prevalence no longer is a cause of excluding a rare disease from a differential diagnosis list, giving this disease a chance to become a final diagnosis, based on merit of supporting clinical data.

- *Disregarding subjective qualities of clinical data*, which are variable and unreliable, remarkably simplifies diagnostic processing without losing accuracy.
- We believe that the *PP value best indicates how strongly a clinical datum supports a diagnosis and more accurately* than does specificity, true positive value, estimated evoking strength [3][7], or any other index attached to a clinical datum; its value does not change unless the sensitivities of clinical data are changed.
- *PP values can be calculated and included in the knowledge base before the diagnostic program is delivered to the user*, avoiding real time calculation.
- Considering the *greatest of the PP values of clinical data present that support a diagnosis, equal to the P of this diagnosis* would appear superior to arithmetically combining values of several redundant supportive clinical data, thereby inappropriately increasing the P of the diagnosis.
- Our *mini-max procedure* for determining the P of a diagnosis overcomes the deficiencies of the typical Bayes formula. In a novel way, the algorithm combines a modified weight averaging Bayes formula with the mini-max principle, to calculate the P of a diagnosis. Original Bayes formula processes sequentially or simultaneously *multiple* clinical data present and absent. Because these clinical data are interrelated, this application violates the independence and incompatibility principles, leading to inaccuracies of calculated P. We apply Bayes averaged formula only to a *single* clinical datum present and a *single* clinical datum absent, in each clinical data pair, which does not violate those principles, and we eliminate prior probability of diseases from it. Once the partial P conferred to the diagnoses by each of all possible clinical data pairs are calculated with the modified averaging Bayes formula, they are integrated into a total P by the mini-max procedure. In this manner, we obtain the benefit of two worlds: original Bayes formula is inaccurate for multiple clinical data processing, but accurate for single pairs of one clinical datum present and one clinical datum absent. Mini-max principle alone does not take into account the appropriate proportional significance of S that is considered by the weight averaged Bayes formula; however, it is appropriate to integrate the partial P values calculated by this formula into a total P of the diagnosis.
- At each step of the diagnostic inquiry, we apply a novel method to select and recommend the *best cost-benefit clinical datum next to investigate*, which heuristically reduces the number of new clinical data searched, thereby considerably shortening the diagnostic process. Unlike other diagnostic programs, our algorithm considers greatest PP value, S, and cost in selecting the best cost-benefit clinical datum.
- *Cost has high priority* in our program and refers not only to the dollar *price*, but also to *discomfort* and *risk*; the maximum of these qualitative levels represents the overall cost.
- *Simultaneously recommending several best cost-benefit clinical data* provides a novel and important advantage over recommending only one at each diagnostic step.
- Ability of our mini-max procedure to *distinguish competitive diagnoses from concurrent diagnoses*.
- Conclusion of diagnostic quest is treated in a manner similar to that used by other authors. However, we have not found literature references to the creation of *confirmation* and *deletion thresholds*.
- The *entire diagnostic process is partitioned into two distinct steps*: the first step uses *probabilities* to obtain final diagnoses; the second step uses *categorical* tools to preclude overlooking related clinical entities and to establish unrelated or related concurrency among diagnosed clinical entities, integrating them into complex clinical presentations. This partitioning eliminates the computational complexity and even impossibility of managing the entire diagnostic process with probabilistic calculations.

- *Risk identifiers, interaction identifiers, and other safety checks*—to preclude overlooking important diagnoses—as well as *empirical treatment, diagnosis by exclusion, and deferred diagnosis* are also important heuristic aspects that complete the benefits of our algorithm.
- The knowledge base must be integrated with all known disease models, including sensitivity (PP values are automatically calculated) and cost category of each clinical datum, risk and interaction identifiers, empirical values for confirmation and deletion thresholds, and all known complex clinical presentation models. When clinical data present and absent are provided, the algorithm is expected to return accurate final diagnoses and complex clinical presentations in an almost automatic manner.

Our diagnostic program could be expanded to include prognostic and therapeutic guidelines. A more ambitious project might be to determine whether other inexact disciplines such as law, sociology, politics, defense, or corporate strategy could benefit from some steps of our algorithm.

CONCLUSIONS

Our algorithm, although somewhat complex, is straightforward, especially when compared to other attempts in this field. It emulates a clinician's diagnostic reasoning. It is logical and mathematically simple. Bayes formula is used with modifications, because it is unable to process properly interdependent clinical data (as are most symptoms) and concurrent diseases. To facilitate implementation and updating of the algorithm, we tend to avoid complicated tools of artificial intelligence, such as causal, hierarchical, and probabilistic trees and networks. The algorithm freely uses heuristic procedures, so as to preclude excessive proliferation of clinical data and diagnoses. It promises to be user friendly because it is expressed in natural language, is rational, and readily understandable. It also is readily translatable to a computer program. Determination of accurate sensitivity of clinical data and integration of clinical entities into clinical presentation models will be labor-intensive. As soon as a prototype program proves the concept viable, a complete knowledge base with all known diseases, clinical data, clinical presentations, and other information can be created. This major task will require a dedicated team. Establishing the sensitivity of all clinical data for all known diseases will be especially time-consuming. This could be accomplished by (a) retrospectively incorporating information from past medical records and publications; (b) prospectively incorporating information from future medical records. The latter task would be an ongoing project, but would be facilitated by universal electronic storage and processing of medical records; (c) basing the sensitivities on the personal estimation of clinicians and specialists, but this is subjective and inaccurate. The expected great benefits of such a medical resource are worth the effort. Perhaps it could be compared with the team effort invested by the creators of many existing multi-volume dictionaries or encyclopedias, or the implementation of some extensive information systems over the Internet.

REFERENCES

[1] LEDLEY RS and LUSTED LB. Reasoning Foundations of Medical Diagnosis. Science, 130 (9): 9-21, July 3, 1959

[2] HENRION M, PRADHAN M, DEL FAVERO B, HUANG K, and O'RORKE P: Why is diagnosis using belief networks insensitive to imprecision in probabilities? Twelfth Conference on Uncertainty in Artificial Intelligence, Portland, OR, 446-454. 1996. SMI-96-0637

[3] MYERS JD, POPLE HE, and MILLER RA. INTERNIST: Can Artificial Intelligence Help? In: Connelly, Benson, Burke, Fenderson, eds. Clinical Decisions and Laboratory Use. Minneapolis: University of Minnesota Press, 1982: 251-269

[4] LUSTED LB. Introduction to Medical Decision Making. Springfield, Illinois, Charles C Thomas, 1968

[5] LUSTED LB. Twenty Years of Medical Decision Making Studies. CH1480-3/79/0000-0004$00.75. 1979 IEEE

[6] FEDER C. Computerized Medical Diagnosis: A Novel Solution to an Old Problem. Infinity Publishing, West Conshohocken, Pennsylvania, 2006

[7] POPLE HE. Heuristic Methods for Imposing Structure on Ill-structured Problems: The Structuring of Medical Diagnosis. In: Szolovits P, ed. Artificial Intelligence in Medicine, AAAS Symposium Series, Boulder, Colorado: West-view Press, 1982: 119-185

[8] MARTIN J: Computer Data-Base Organization. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1977

[9] CHAKRAVARTY S and SHAHAR Y. A Constraint-Based Specification of Periodic Patterns in Time-Oriented Data. Sixth International Workshop on Temporal Representation and Reasoning (TIME-99), Orlando, FL, 29-40. 1999. SMI-1999-0766

ALTMAN RB. AI in Medicine: The Spectrum of Challenges from Managed Care to Molecular Medicine. AI Magazine 20(3):67-77, 1999. SMI-1999-0770

BLEICH HL. Computer-Based Consultation: Electrolyte and Acid-Base Disorders. The American Journal of Medicine 53: 285-291, November 1972

BLOIS MS, TUTTLE MS, and SHERERTZ DD. RECONSIDER: A Program for Generating Differential Diagnoses. IEEE: 263-268, 1981

CHAKRAVARTY S and SHAHAR Y. Acquisition and Analysis of Periodic Patterns in Time-Oriented Clinical Data. 2000. SMI-2000-0822

DE DOMBAL FT, LEAPER DJ, STANILAND JR, McCANN AP, and HORROCKS JC. Computer-Aided Diagnosis of Acute Abdominal Pain. British Medical Journal 2: 9-13, 1972

ESHELMAN L, EHRET D, McDERMOTT JP, and TAN M (1987.) MOLE: A Tenacious Knowledge Acquisition Tool. International Journal of Man-Machine Studies, 26: 41-54

GENNARI J, MUSEN MA, FERGERSON RW, GROSSO WE, CRUBEZY M, ERIKSSON H, NOY NF, TU SW: The Evolution of Protégé: An Environment for Knowledge-Based Systems Development. 2002. SMI-2002-0943

GINI RA and FEDER C. Informatica Clinica: Presente y Futuro. La Semana Medica (Argentina) 157: 113-128, 1980

GORRY GA, KASSIRER JP, ESSIG A, and SCHWARTZ WB. Decision Analysis as the Basis for Computer-Aided Management of Acute Renal Failure. The American Journal of Medicine 55: 473-484, 1973

GORRY GA, PAUKER SG, and SCHWARTZ WB. The Diagnostic Importance of the Normal Finding. The New England Journal of Medicine 486-489, March 2, 1978

GREENS RA, PELEG M, BOSWALA AA, TU S, PATEL VL, and SHORTLIFFE EH. Sharable Computer-Based Clinical Practice Guidelines: Rationale, Obstacles, Approaches, and Prospects. Medinfo, London, UK, 2001. SMI-2001-0860

HUANG K, HENRION M: Efficient Search-Based Inference for Noisy-OR Belief Networks: Top Epsilon. Proceedings of the Twelfth Conference of Uncertainty in Artificial Intelligence, 325-331. Aug 1996, Portland, OR. SMI-96-0640

ILIAD® 4.5 Diagnostic and Reference Tool for Physicians and Medical Professionals. User Guide. 1998

JAAKKOLA TS, JORDAN MI: Variational Probabilistic Inference and the QMR-DT Network. Sun May 9, 16:22:01 PDT 1999

LUDWIG DW. INFERNET – A Computer-Based System for Modeling Medical Knowledge and Clinical Inference. Proceedings of the Fifth Annual Symposium on Computer Applications in Medical Care: 243-249, November 1981

MIDDLETON B, SHWE M, HECKERMAN D, HENRION M, HORVITZ E, LEHMANN H, & COOPER G: Probabilistic Diagnosis using a reformulation of the INTERNIST-1/QMR Knowledge Base II. Evaluation of Diagnostic Performance. Section on Medical Informatics Technical report SMI-90-0329, Stanford University, 1990

MILLER RA, POPLE HE, and MYERS JD. INTERNIST-I, An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine. The New England Journal of Medicine, 1982: 468-476

MUSSEN MA, GENNARI JH, and WONG WW: A Rational Reconstruction of INTERNIST-I Using PROTÉGÉ-II. Nineteenth Annual Symposium on Computer Applications in Medial Care. New Orleans, LA, 289-293. 1995. SMI-95-0574

MUSEN MA: Modeling for Decision Support. Section on Medical Informatics. Stanford University School of Medicine. Stanford, CA 94305-5479. SMI-98-0739

NOY NF and MUSEN MA. Ontology Versioning as an Element of an Ontology-Management Framework. Stanford Medical Informatics, Stanford University, 251 Campus Drive, Stanford, CA 94305, USA. SMI-2003-0961. March 31, 2003

PATRICK EA: Decision Analysis in Medicine: Methods and Applications. CRC Press, West Palm Beach, FL, 1979

PERLROTH MG, and WEILAND DJ. Fifty Diseases: Fifty Diagnoses. Year Book Medical Publishers, 1981

POPLE HE, MYERS JD, and MILLER RA: Dialog: A Model of Diagnostic Logic for Internal Medicine. Fourth International Joint Conference on Artificial Intelligence. Tbilisi, Georgia, URRS, 3-8 September 1975, Volume Two. 1975: 848-855

PRADHAN M, DAGUM P: Optimal Monte Carlo Estimation of Belief Network Inference. Proceedings of the Twelfth Conference of Uncertainty in Artificial Intelligence, 446-453. Aug 1996, Portland, OR. SMI-96-0638

SHORTLIFE EH. Computer-Based Medical Consultations: MYCIN. American Elsevier Publishing Company, 1976

SHORTLIFFE EH. The Next Generation Internet and Health Care: A Civics Lesson for the Informatics Community. In C.G. Chute, Ed., 1998 AMIA Annual Symposium, Orlando, FL, 8-14. 1998. SMI-98-0730

SHWE M, MIDDLETON B, HECKERMAN D, HENRION M, HORVITZ E, LEHMANN H, & COOPER G. Probabilistic Diagnosis Using a Reformulation of the INTERNIST-1/QMR Knowledge Base I. The Probabilistic Model and Inference Algorithms. Methods of Information in Medicine, 30(4):241-255, 1991. SMI-90-0296

SZOLOVITS P and PAUKER SG. Categorical and Probabilistic Reasoning in Medical Diagnosis. Artificial Intelligence. 11: 115-144, 1978

TU SW et al: Modeling Guidelines for Integration into Clinical Workflow. Department of Medicine, Stanford University School of Medicine, Stanford, CA, USA. SMI-2003-0970

VAN BEMMEL JH, MUSSEN MA (eds): Medical Informatics. Springer, 1997

WEINSTEIN MC and FINEBERG HV: Clinical Decision Analysis. W. B. Saunders Company, 1980

WEISS S, KULIKOWSKI CA, and SAFIR A. Glaucoma Consultation by Computer. Comput. Biol. Med. 8: 25-40, 1978