

Zero Order and Nonzero Order Decision Rules in Medical Diagnosis

Abstract: In searching for the optimal solution to the medical diagnostic problem, it seems useful to distinguish between different possible decision rules (strategies). Two different classes of decision rules are considered: nonzero order decision rules and constant or zero order decision rules. For each class, solution methods as well as heuristic approaches to finding the optimal member of the class are discussed. As an exercise, the diagnosis of a hematologic disease belonging to the group macrocytic anemia is considered.

Introduction

Decision rules [1-7] arise, for example, in the study of statistical games and in the context of various multistage stochastic programming models. A decision rule may be defined as a vector-valued function that maps a random variable into a decision. According to a nonzero order decision rule in a multistage stochastic programming model, the values of the decision variables are based upon values of random variables that are observed during the time horizon; the values of the decision variables for stage t are specified as explicit functions of the outcomes of the random variables for the stages $j = 1, \dots, t-1$. A special subclass of the general class of nonzero order rules is formed by the so-called constant or zero order decision rules. According to a zero order rule the values of all decision variables have to be determined at the inception of the time horizon. Thus, the main difference between the two is that according to the nonzero order decision rule the exact value of a decision variable with respect to stage t can only be computed after the outcomes of all random variables concerning the preceding $t-1$ stages have been observed; according to the zero order one the value of the decision variable is known exactly at time $t = 0$, the inception of the period of planning.

If the search for the optimal decision rule is restricted to a subclass of the nonzero order decision rules, the overall optimal decision rule might not be found. In some cases, the restricted problem may have no feasible solution at all. There may be good reasons to restrict the search to the zero order rules, however. For example, in multiperiod regional planning problems, a solution that specifies the level of execution of a decision variable of some stage t , $t > 1$, conditional upon the outcomes of the random variables of preceding stages, generally has no practical meaning, because preparations for execution on that level have to be made at the inception of the time horizon.

The medical diagnostic problem for the physician consists globally of two phases:

1. The development of a diagnosis from a given set of clinical signs (the set of observed test outcomes).
2. The determination of an appropriate sequence of diagnostic tests to perform on the patient in order to collect further information, if needed.

The development of the experimental computer program HEME illustrates the use of computer technology to assist in the diagnostic activity [8, 9]. The program accepts as input information about a patient (the set of observed test outcomes) and provides as output an ordered list of suggested diagnoses, an analysis of the logic behind these diagnoses, and a list of tests relevant to these diagnoses and not yet performed.

Solving the second phase of the medical diagnostic problem is a prerequisite for the development of such a program. This phase can be seen to fit in the class of stochastic programming problems or, more precisely, in the class of multistage stochastic programming models. In a number of respects it differs fundamentally from the models that are usually dealt with in the literature, however. In most stochastic models the decision variables are continuous, and it is supposed that their values do not affect the probability distribution of the random variables. However, the decision variables of the model for solving the second phase of the medical diagnostic problem are zero-one variables by nature, representing yes or no decisions such as: "whether or not to perform some test" and "whether or not some disease is definitely being diagnosed," and their values clearly affect the probability distribution of the random variables (possible test outcomes).

This paper aims at interpreting nonzero order and zero order decision rules within the special context of solving the second phase of the medical diagnostic problem. In

Table 1 Notation.

Notation	Meaning
D	set of possible diseases
T	set of available tests
d_i	i th disease
T_j	j th test
e_{jq}	q th possible outcome of T_j
e_q	observed outcome of q th test of a test sequence performed on patient
$p(e_{jq} d_i)$	conditional probability of obtaining e_{jq} if patient has d_i
$p(d_i e_1, \dots, e_n)$	probability that patient has d_i after $e_1, \dots, e_n$ have been observed.
$p(d_i)$	a priori probability that patient has d_i
$P(D e_1, \dots, e_n)$	vector of updated probabilities of diseases after $e_1, \dots, e_n$ have been observed
L	loss matrix, in which element l_{ij} in i th row and j th column denotes costs (loss) of diagnosing d_i when patient actually has d_j
L_i	i th row of L
$P(e_{jq} D)$	vector consisting of elements $p(e_{jq} d_1), \dots, p(e_{jq} d_n)$
c_j	costs of performing T_j
t'_j	waiting time for outcome of T_j
t''_j	execution time of T_j
$t_j = t'_j + t''_j$	total processing time of T_j
$C(e_1, \dots, e_n)$	expected costs of "best" terminal decision (diagnosis) after $e_1, \dots, e_n$ have been observed
$p(e_{jq} e_1, \dots, e_n)$	probability of obtaining e_{jq} given that $e_1, \dots, e_n$ have been observed
N	number of different diseases
M	number of different tests available
$n(j)$	number of possible outcomes of T_j

considering a particular (ordered) test sequence, a physician should weigh the expected value of the test results against the costs of the tests. On the one hand he seeks to minimize the consequences of possible misdiagnosis, while on the other hand he aims at keeping the number of costly diagnostic tests as low as possible. In searching for the optimal solution to this problem, different starting points are possible, generally leading to different solutions [10-14]. Such different starting points can be interpreted as different decision rules within the framework of stochastic programming.

For both nonzero and zero order decision rules, algorithms to determine the optimal member are discussed, as well as heuristic approaches to find a good suboptimal

solution. An experimental computer program implementing the algorithms to determine the optimal nonzero and the optimal zero order rules has been written in APL and tested on a computer. An exercise is presented in which this program is used for hematologic disease diagnoses derived from HEME [8]. In presenting the exercise the practical problems that arise in specifying the loss for misdiagnosis and the costs of performing tests are discussed.

Notation and assumptions in analyzing the medical diagnostic problem

The notation used throughout this paper is presented in Table 1. In analyzing the medical diagnostic problem, usually a number of rather restrictive assumptions are supposed to be valid. Often a list of all underlying assumptions is lacking, however, which sometimes makes it difficult to understand further analysis and to evaluate the degree of artificiality of the model involved. We therefore give an explicit survey of all assumptions that are supposed to be valid in further analysis.

1. The set D of possible diseases is known and fixed, and all diseases are different.
2. The set T of available tests is known and fixed, and all tests are different. Each test is characterized by a number of possible outcomes.
3. All tests remain available for execution during an unlimited time period.
4. The patient has one and only one of the elements of D .
5. Each pair of diseases differs in at least one attribute (symptom), the presence or absence of which can be revealed by at least one of the available tests.
6. The a priori probability distribution of the diseases is known.
7. Of each test the conditional probability of obtaining some outcome, given that the patient has some disease, is known for each of the possible diseases and outcomes of the test.
8. For each ordered pair of diseases i and j the costs of misdiagnosing disease j rather than disease i as existing in the patient are known.
9. The total processing time of each test (consisting of its execution time and the waiting time for its outcome) and the associated costs of its performance are fixed.
10. The outcome of a test can be determined with certainty.
11. The probability of obtaining some test outcome given that the patient has some disease is independent of the sequence of tests already performed on the patient.

12. Both the costs of misdiagnosis and the costs of performing tests are given in the same units.
13. Treatment of the patient is started after some disease has been diagnosed.

Some of the constraints are more restrictive than others and have further reaching theoretical consequences as a result. Assumption 11 deserves special attention. It implies that $p(e_{jq}|d_i, e_1, \dots, e_n) = p(e_{jq}|d_i)$ holds for all j, q , and i . This means that no distinction has to be made among sequences of the same tests but performed in a different order. In the next section we state two properties that are based upon the validity of assumption 11. According to assumptions 2 and 11 Bayes' theorem can be applied directly for updating the a priori probabilities of the diseases after some sequence of tests has been performed and the corresponding outcomes have been observed. According to Bayes' theorem, the updated probability of disease d_i after the first test outcome e_1 has been observed can be computed by

$$p(d_i|e_1) = \frac{p(d_i, e_1)}{p(e_1)} \\ = \frac{p(d_i) \cdot p(e_1|d_i)}{p(d_1) \cdot p(e_1|d_1) + \dots + p(d_N) \cdot p(e_1|d_N)} \quad (1)$$

However, had $d_1, \dots, d_N$ not formed a set of mutually exclusive events as required by assumption 4, $p(e_1)$ would have to be computed by

$$p(e_1) = \sum_{i=1}^N p(d_i, e_1) - \sum_{\substack{i,k \\ i < k}} p(d_i, d_k, e_1) \\ + \sum_{\substack{i,k,z \\ i < k < z}} p(d_i, d_k, d_z, e_1) - \dots + p(d_1, \dots, d_N, e_1).$$

Otherwise the computed value of $p(d_i|e_1)$ will generally be an underestimation of the true value.

Let us now turn to the optimization criterion. In most of the literature on sequential decision making in medical diagnosis, the optimization criterion consists of minimizing the total expected costs—the so-called “E-model” in the literature on stochastic programming [7]. Of course, other criteria can easily be imagined, reflecting other attitudes of the decision maker towards risk, e.g., minimization of the expected value of the variance of the total costs. According to that criterion the goal of the decision maker is to minimize the risk of obtaining extreme final results. Furthermore, time-based criteria such as minimization of the expected total time of examination may be the objective. Finally, one can think of minimizing more than one criterion at the same time. In the main part of our analysis, the optimization criterion

is the usual one of minimizing the total expected costs. However, sometimes mixtures of cost-based and time-based criteria are discussed.

Basic theory

According to assumptions 6, 7, 8, and 9, stated in the previous section, the decision maker is supplied with all the necessary data to reach a rational decision. To construct a current view of the problem after some set of outcomes has been obtained, access to the a priori probabilities of the diseases, as well as to the conditional probabilities of each of the possible test outcomes given each of the diseases, is a prerequisite. Furthermore, one requires access to the data of the loss matrix without which it would be impossible to evaluate the cost consequences of stopping further testing and diagnosing some disease on the basis of the accumulated information gathered at some stage of the decision making process. In a later section, in discussing our test case, the practical problems in specifying all these relevant data (especially of specifying $c_j, j = 1, \dots, M$, and the elements of L , all in the same units) are emphasized.

As discussed earlier, the underlying assumptions imply that Bayes' theorem can be applied directly to update the a priori probabilities of the diseases after some sequence of tests has been performed and the corresponding outcomes have been observed. Bayes' rule has been discussed rather extensively in most literature concerning the inference aspects of medical diagnosis, and a number of motives are advanced in favor of the use of this rule in this special context [10–14]. Equation (1) can be easily generalized for n findings in the following way. Let R denote the set of all possible permutations of n findings $e_1, \dots, e_n$. So, R consists of $n!$ elements $R_1, R_2, \dots, R_n$, each element representing a different permutation of $e_1, \dots, e_n$. Then the following property can be formulated.

Property 1 If assumption 11 is valid, then

$$p(d_i|R_j) = p(d_i|R_k), \quad \forall R_j, R_k \in R; i = 1, \dots, N.$$

Thus, property 1 states that $p(d_i|e_1, \dots, e_n)$ does not depend on the order in which the test results $e_1, \dots, e_n$ are obtained.

From property 1 and from the observation that the probability of obtaining some test outcome e_{jq} is completely determined by the current, updated probability distribution of the diseases and the conditional probabilities $p(e_{jq}|d_i), i = 1, \dots, N$, the next property follows directly.

Property 2 If assumption 11 is valid, the probability of obtaining some test outcome $e_{jq} \notin \{e_1, \dots, e_n\}$ from test T_j after having observed the test results $e_1, \dots, e_n$ does not depend on the order in which those results have been obtained;

Table 2 Decision losses.

Diagnosis/actual disease	$d_1 \cdots d_N$
d_1	$l_{1,1} \cdots l_{1N}$
$\vdots$	$\vdots \cdots \vdots$
d_N	$l_{N1} \cdots l_{NN}$

$$p(e_{jq}|R_j) = p(e_{jq}|R_k), \quad \forall R_j, R_k \in R.$$

Let us now turn to the loss matrix L . The matrix (Table 2) contains the data of the decision losses.

The elements l_{ii} , $i = 1, \dots, N$, on the main diagonal of L are all equal to 0, because no costs are associated with diagnosing a disease that the patient actually has. The elements l_{ij} , $i = 1, \dots, N$; $j = 1, \dots, N$; $i \neq j$, correspond to the loss of treating the patient in the wrong way and/or failing to treat the patient in the right way, if possible. In general, L will be asymmetric. The expected losses (costs) of diagnosing d_i after $e_1, \dots, e_n$ have been observed, denoted by $c(d_i, e_1, \dots, e_n)$, are found by computing

$$c(d_i, e_1, \dots, e_n) = P(D|e_1, \dots, e_n)L_i, \quad (3)$$

where $P(D|e_1, \dots, e_n)$ and L_i are defined in Table 1.

Now the expected costs of the "best" terminal decision, given that $e_1, \dots, e_n$ have been observed, can be defined.

Definition The expected costs of the best terminal decision (diagnosis) after $e_1, \dots, e_n$ have been observed, denoted by $\bar{C}(e_1, \dots, e_n)$, are defined by

$$\bar{C}(e_1, \dots, e_n) = \min_{i=1, \dots, N} \{c(d_i, e_1, \dots, e_n)\}. \quad (4)$$

In general the disease diagnosed according to the best terminal decision will not be unique.

Nonzero order and zero order decision rules

As discussed earlier one can think of the medical diagnostic problem as a specific multistage stochastic programming problem. If we put the problem in this framework we can distinguish more clearly between the different starting points that are possible in searching for an optimal solution by considering them as different decision rules. We now consider the interpretation of a nonzero order and a zero order decision rule in the context of the medical diagnostic model.

In a nonzero order decision rule the k th stage decision is made conditional upon the accumulated information after performing the $k - 1$ preceding tests. Thus, only with respect to the first decision of the sequence does there exist absolute certainty at the inception time t_0 .

In a zero order rule absolute certainty about the performance of all tests exists at the inception time t_0 , and there is no need to wait for the outcome of some test before performing the next one. It is likely that the decision maker (physician) is more interested in a solution of this type. There is no restriction on the order in which the tests have to be executed, and the total examination time will be much shorter on the average. Moreover the preparation for the execution of all tests of the sequence can be started at t_0 . Preparation may involve readying the patient for the test, scheduling needed activities for the performance of the test (e.g., reserving the operating room), etc.

In the following sections a general solution method as well as heuristic approaches for finding the optimal nonzero order rule and zero order rule are described.

• Determination of the optimal nonzero order decision rule

Usually the sequential decision making problem is represented by a so-called decision tree, consisting of a set of nodes that are interconnected in space by a system of directed curves (branches). Two types of nodes can be distinguished: *decision nodes* and *artificial nodes*. The decision nodes correspond to points at which a number of test results have been observed and at which the decision has to be made whether to continue testing or to stop testing and diagnose a disease of which the total expected costs of misdiagnosis are minimal. Each of the branches emanating from some decision node n_k leads to an artificial node that corresponds to the selection of some test T_j to perform. Each of the branches emanating from some artificial node leads to a decision node that corresponds to one of the possible outcomes of the test involved. In Fig. 1 the decision tree has been laid out for a small example of two tests being available and each of the two tests having two different possible outcomes. The artificial nodes are marked by θ .

In describing how to find the optimal nonzero order decision rule, the following definitions and notations are useful.

Notation The set of decision rules that spring from branching from node n_k , called the descendants of n_k , is denoted by $I(n_k)$.

Definition Let node n_k be branched from by selecting test T_j . Then the decision nodes corresponding to the different possible outcomes of T_j are called the *direct descendants* of n_k with respect to T_j .

Notation The expected costs of the optimal decision rule of node n_k are denoted by F_k .

Notation Let n_k correspond to the observation of the test results $e_1, \dots, e_n$. Then $p(d_i|n_k)$, $P(D|n_k)$, $\bar{C}(n_k)$, and

$p(e_{jq}|n_k)$ correspond to $p(d_i|e_1, \dots, e_n)$, $P(D|e_1, \dots, e_n)$, $\bar{C}(e_1, \dots, e_n)$, and $p(e_{jq}|e_1, \dots, e_n)$, respectively.

Definition Node n_k is said to belong to level q of the decision tree if n_k has been reached after performing q tests.

Let n_k be any node of the decision tree that has been reached after performing r tests, e.g., $T_1, T_2, \dots, T_r$, and corresponding to the observation of the test results $e_1, \dots, e_r$. Let us further assume that the optimal decision rules of all descendants of n_k have already been determined. Given the updated probability distribution $P(D|n_k)$ at node n_k , the probability of obtaining test result e_{jq} after performing test $T_j, j \notin \{1, \dots, r\}$, can be computed by

$$p(e_{jq}|n_k) = p(d_1|n_k) \cdot p(e_{jq}|d_1) + \dots + p(d_N|n_k) \cdot p(e_{jq}|d_N). \quad (5)$$

Let $n_{k+1}, n_{k+2}, \dots, n_{k+n(j)}$ be the direct descendants of n_k with respect to test T_j . As $F_{k+1}, F_{k+2}, \dots, F_{k+n(j)}$ are already known by assumption, the expected costs of continuing testing with test T_j , denoted by $\bar{C}(T_j, n_k)$, can be computed by

$$\bar{C}(T_j, n_k) = c_j + \sum_{q=1}^{n(j)} p(e_{jq}|n_k) F_{k+q}, \quad (6)$$

where c_j denotes the costs of performing test T_j . The expected costs of the best terminal decision at node n_k , denoted by $\bar{C}(n_k)$, are found by

$$\bar{C}(n_k) = \min_{i=1, \dots, N} \{P(D|n_k) L_i\}. \quad (7)$$

Finally, the optimal decision rule of n_k can be determined by taking

$$\min \{\bar{C}(n_k), \bar{C}(T_1, n_k), \bar{C}(T_2, n_k), \dots, \bar{C}(T_M, n_k)\}. \quad (8)$$

At this point two observations can be made. First, the expected costs of the optimal decision rule of all members of $I(n_k)$ have to be computed before F_k can be computed. Secondly, the optimal decision rule at node n_k can be found independently of those of the nodes that do not belong to $I(n_k)$. From this it follows that the optimal decision rule of node n_0 , the ultimate goal, can be found by a procedure of averaging out and folding back the decision tree. To start with, the expected costs of the optimal decision rule of all nodes that belong to level M of the decision tree are determined. At these nodes there is no other alternative than accepting the best terminal decision. Next the expected costs of the optimal decision rules of all nodes that belong to the $(M-1)$ tree level are determined. The choice at each of these nodes is restricted to two alternatives: "Proceed testing by performing the test left," and "stop testing and accept the




Figure 1 Tree for example.

best terminal decision given the information gathered so far." By tracking further backwards in this way, we finally find, at least theoretically, the optimal decision rule of n_0 .

Notice that the optimal decision rule of n_0 generally will have a rather complex structure, such as, "Perform as the first test T_i ; if the outcome e_{i1} is observed, then stop further testing and diagnose disease d_j ; if e_{i2} is observed, then perform T_k as the next test," and so on. Nodes marked with the optimal decision rule "stop further testing" may occur at any level of the tree. They explain why two test sequences (two decision rules), which consist of the same tests but ordered in a different way, may have different associated expected costs, though the probability of obtaining some set of test outcomes is the same for both sequences. This last observation is crucial and is illustrated by the following small example.

We distinguish three diseases, d_1, d_2, d_3 , and two tests, T_1 and T_2 . Each test has two possible different outcomes. Further, the data in Table 3 are known. The resulting decision tree has the shape shown in Fig. 1.

First, the function values F_5, F_6, F_7 , and F_8 of the nodes on the lowest level of the tree can be computed.

$$F_5 = \bar{C}(n_5) = \min_{i=1,2,3} \{P(D|e_{1,1}, e_{2,1}) \cdot L_i\} = 80, \text{ for } i = 3,$$

$$F_6 = \bar{C}(n_6) = \min_{i=1,2,3} \{P(D|e_{1,1}, e_{2,2}) \cdot L_i\} = 259, \text{ for } i = 1,$$

$$F_7 = \bar{C}(n_7) = \min_{i=1,2,3} \{P(D|e_{1,2}, e_{2,1}) \cdot L_i\} = 80, \text{ for } i = 3,$$

$$F_8 = \bar{C}(n_8) = \min_{i=1,2,3} \{P(D|e_{1,2}, e_{2,2}) \cdot L_i\} = 259, \text{ for } i = 2.$$

Subsequently, the expected cost of the optimal decision rules of the nodes on the first level of the tree can be determined. To this purpose $p(e_{2,1}|e_{1,1})$, $p(e_{2,2}|e_{1,1})$, $p(e_{2,1}|e_{1,2})$, $p(e_{2,2}|e_{1,2})$, $p(e_{1,1}|e_{2,1})$, $p(e_{1,2}|e_{2,1})$, $p(e_{1,1}|e_{2,2})$, and $p(e_{1,2}|e_{2,2})$ have to be computed first. We find

Table 3 Conditional probabilities, loss matrix, a priori probabilities, and costs.

Conditional probabilities				
disease/test outcomes	$e_{1,1}$	$e_{1,2}$	$e_{2,1}$	$e_{2,2}$
d_1	0.95	0.05	0.05	0.95
d_2	0.05	0.95	0.05	0.95
d_3	0.50	0.50	0.80	0.20

Loss matrix			
diagnosis/actual disease	d_1	d_2	d_3
d_1	0	500	1000
d_2	500	0	1000
d_3	2000	2000	0

A priori probabilities of diseases	
disease	a priori probability
d_1	0.2
d_2	0.2
d_3	0.6

Costs associated with performing tests	
test	costs
T_1	200
T_2	200

$F_1 = \min \{ \min_{i=1,2,3} \{ P(D|e_{1,i}) \cdot L_i \}, c_2 + p(e_{2,1}|e_{1,1}) \cdot F_5 + p(e_{2,2}|e_{1,1}) \cdot F_6 \} = 369.5$, whereas the optimal decision rule of n_1 reads: "Perform test T_2 ; if $e_{2,1}$ is observed, then diagnose d_3 ; if $e_{2,2}$ is observed, then diagnose d_1 ."

$F_2 = 80$, whereas the optimal decision rule of n_2 reads: "Stop further testing and diagnose disease d_3 ."

$F_3 = 369.5$, whereas the optimal decision rule 3 of n_3 reads: "Perform test T_2 ; if $e_{2,1}$ is observed, then diagnose d_3 ; if $e_{2,2}$ is observed, then diagnose d_2 ."

$F_4 = 430$, whereas the optimal decision rule of n_4 reads: "Stop further testing and diagnose d_1 or d_2 ."

Besides F_1 , F_2 , F_3 , and F_4 , we need to know $p(e_{1,1})$, $p(e_{1,2})$, $p(e_{2,1})$, $p(e_{2,2})$, and $\bar{C}(n_0)$ in order to determine the optimal decision rule of n_0 . We have

$$\bar{C}(n_0) = \min_{i=1,2,3} \{ P(D) \cdot L_i \} = 700 \text{ for } i = 1, 2, \text{ and}$$

$$p(e_{1,1}) = p(e_{1,2}) = p(e_{2,1}) = p(e_{2,2}) = 0.5.$$

Consequently,

$$\bar{C}(T_1, n_0) = c_1 + p(e_{1,1}) \cdot F_1 + p(e_{1,2}) \cdot F_3 = 569.5, \text{ and}$$

$$\bar{C}(T_2, n_0) = c_2 + p(e_{2,1}) \cdot F_2 + p(e_{2,2}) \cdot F_4 = 455.$$

The optimal decision rule of n_0 follows from determining $\min \{ \bar{C}(n_0), \bar{C}(T_1, n_0), \bar{C}(T_2, n_0) \}$,

and turns out to be performing T_2 as the first test. The complete decision rule of n_0 reads: "Perform T_2 as the first test; stop further testing if the outcome $e_{2,1}$ is observed, and diagnose d_3 ; if $e_{2,2}$ is observed, then stop further testing and diagnose d_1 or d_2 ."

Let us look in somewhat more detail at the solution arrived at. If T_2 is selected as the first test to perform, there exists absolute certainty that no further test has to be performed. However, if T_1 is selected as the first test, T_2 always has to be performed as the next one, regardless of the outcome of T_1 . If no decision nodes with optimal decision rule: "Stop further testing" occurred in the tree except at those nodes at the lowest level, then both decision rules at node n_0 : "Start with T_1 " and "Start with T_2 " would have had the same associated expected costs.

This simple example illustrates clearly that different ordered sequences of the same tests may involve different expected costs. One of the main reasons for this is that the discrimination power of a test (the ability of a test to rule out and/or to support strongly the presence of a number of diseases) generally differs, even strongly differs, from node to node as each node corresponds to accumulated information of a different type. Some tests may be quite well suited for separating the members of some specific subset of diseases, and therefore performing such a test only makes sense at that stage of the decision making process where the probability of the presence of that subset is sufficiently high. Other tests may be much more suitable to distinguish between different groups of diseases but have no specific discrimination power to single out some particular disease. Also, the costs of misdiagnosis and the costs associated with the performance of each of the tests clearly play an important role.

Heuristic approaches for finding the optimal nonzero order decision rule

If the number of available tests is very large and/or there are a number of tests with a large number of relevant outcomes, the general solution method, as described in the last section, becomes very time consuming, if not practically impossible to execute. In this case one has to resort to heuristic solution methods that yield good sub-optimal solutions but not necessarily optimal ones.

One general way to limit the growth of the decision tree is to put extra restrictions on the solution to the problem by, for example, requiring that the total expected time for examination of the patient should not exceed some specified level. Then the search for the optimal nonzero order rule remains restricted to the class of rules that can be applied when the total expected exam-

ination time is less than the given upper bound. Another general method is to restrict the depth of the analysis at each node of the tree. For example, Gorry and Barnett consider a depth of analysis of one at each node [10]. However, this approach generally gives rise to decision rules according to which further testing is stopped prematurely, as the function values of the direct descendants of some node n_k are systematically overestimated by setting them equal to the expected costs of the best terminal action at the node. The optimal decision rule for the example of Gorry and Barnett is to start by performing test T_3 if a depth of analysis of one at each node is considered, whereas the optimal decision rule with an unrestricted depth of analysis would be to start with test T_2 .

We now discuss briefly two possible strategies to select a restricted number of tests out of the set of candidates that are available to branch from some node n_k . Let D_j be a subset of D with the following properties:

1. The outcome of test T_j is characteristic for each member $d \in D_j$. This means that one of the possible outcomes $e_{j1}, \dots, e_{jn(j)}$ of T_j will occur with a probability greater than some user specified level γ if d is true. Thus $d \in D_j$ implies: There exists some r such that $p(e_{jr}|d) > \gamma$.
2. The outcome of T_j is not in the same way characteristic for all elements of D_j .

If γ is chosen close to 1, D_j will generally consist of a relatively small number of elements. We denote the complement of D_j in D by $\bar{D}_j$ and the set of indices of all elements that belong to $\bar{D}_j$ by $\bar{I}_j$. The conditional probability that $\bar{D}_j$ is true, given the information gathered at node n_k , can be computed by

$$p(\bar{D}_j|n_k) = \sum_{i \in \bar{I}_j} p(d_i|n_k). \quad (9)$$

Now, order the candidate tests to branch from n_k in the sequence $T_{r(1)}, \dots, T_{r(n)}$ such that

$$p(\bar{D}_{r(1)}) \leq p(\bar{D}_{r(2)}) \leq \dots \leq p(\bar{D}_{r(n)})$$

holds. One of the following two strategies may be applied to restrict the number of tests to branch from node n_k :

1. Branch from n_k by means of the first q tests $T_{r(1)}, \dots, T_{r(q)}$.
2. Determine q^* such that $p(\bar{D}_{q^*}) \leq \beta$, where β is some user specified level. Branch from n_k by means of the first q^* tests $T_{r(1)}, \dots, T_{r(q^*)}$.

The philosophy of both strategies is to restrict branching from n_k to those tests that are expected to yield the most useful information at that stage of the sequential decision making process. (If the probability of $\bar{D}_j$ with respect to test T_j is large, relatively less valuable information is likely to be obtained by performing T_j at that stage of the sequential decision making process.)

Zero order decision rules

As discussed previously the decision maker can choose between searching for a nonzero order decision rule and searching for a zero order decision rule; in addition mixtures of both strategies can well be imagined. Obviously the expected costs of the optimal zero order rule or any mixture of a zero order rule and a nonzero order one are higher than or equal to the expected costs of the proper optimal nonzero order rule, since the first one can be considered as a particular case of the second one. We now discuss how, at least theoretically, the optimal zero order rule can be found.

Define the optimal test sequences of length r as that sequence of r different tests for which the total expected costs are minimal. The total number of different test sequences of length r , denoted by $M(r)$, equals $\binom{M}{r}$.

Let us denote the possible different test sequences of length r by $S(1, r), \dots, S(M(r), r)$. Further, denote the set of indices of the tests that form part of $S(k, r)$ by $I(k, r)$. Then the number of possible different combinations of test outcomes after the r tests of $S(k, r)$ have been performed, denoted by $q(k, r)$, is

$$q(k, r) = \prod_{i \in I(k, r)} n(i). \quad (10)$$

Each combination of test outcomes can be represented as one of the terminal nodes of a decision tree, say $n_1, n_2, \dots, n_{q(k, r)}$. According to property 2, $p(n_i)$, the probability of reaching node n_i , does not depend on the order in which the r tests are performed, and the same holds for $P(D|n_i)$. (See property 1.) The expected costs of the best terminal action at node i , denoted by $\bar{C}(n_i)$, can be computed by

$$\bar{C}(n_i) = \min_{j=1, \dots, N} \{P(D|n_i) \cdot L_j\}.$$

The total expected costs implied by $S(k, r)$, including not only the costs of performing the tests but also the expected costs of making incorrect diagnoses, denoted by $\bar{C}(S(k, r))$, can be computed successively by

$$\bar{C}(S(k, r)) = \sum_{i=1}^{q(k, r)} p(n_i) \cdot \bar{C}(n_i) + \sum_{i \in I(k, r)} c_i, \quad (11)$$

where c_i denotes the costs of performing T_i .

To find the optimal test sequence of length r , and the costs associated with that sequence, determine j^* such that

$$\bar{C}(S(j^*, r)) = \min_i \{\bar{C}(S(i, r))\}. \quad (12)$$

Then $S(j^*, r)$ is the optimal test sequence of length r , and $\bar{C}(S(j^*, r))$ are the costs associated with that optimal sequence. Let $\bar{C}(S(j^*, 0))$ denote the expected costs of the best terminal action given the a priori probability distribution of the diseases. Then the optimal zero order

Table 4 Diseases of the macrocytic anemia group.

No.	Disease
1	pernicious anemia
2	gastrectomy syndrome
3	intestinal disorders
4	competitive parasite
5	poor nutrition
6	impaired absorption
7	pregnancy
8	hyperactive hematopoiesis
9	skin and neoplastic diseases
10	folic acid antagonists
11	vitamin C
12	not B ₁₂ or F.A.
13	juvenile P.A.

decision rule is found by determining r^* such that

$$\bar{C}(S(j^*, r^*)) = \min_{r=0,1,\dots,M} \{\bar{C}(S(j^*, r))\}. \quad (13)$$

As with the optimal nonzero order rule, the algorithm to find the optimal zero order rule is computationally infeasible if the number of available tests is too large and/or the number of possible outcomes of a number of tests is too large. It is then necessary to resort to heuristic methods to find satisfactory suboptimal solutions. One method is to put a restriction on the total examination time of the patient. This implies that time consuming tests will only occur in test sequences of relatively short length. In addition, use can be made of the a priori knowledge that a combination of mutually "conflicting" test outcomes is not very likely to be obtained. According to this philosophy the probability of reaching some node n_i , corresponding to a combination of conflicting test outcomes, is supposed to be such a small number that the support of its term $p(n_i) \cdot \bar{C}(n_i)$ in (11) can be neglected. Thus $\bar{C}(n_i)$ only has to be computed for a restricted number of terminal nodes of the tree.

Experiments in hematologic disease diagnosis

The algorithms to determine the optimal nonzero and zero order decision rules described earlier have been programmed in APL and tested on a computer. To exercise the program on a real-life case, the diagnosis of a hematologic disease, belonging to the group macrocytic anemia, was considered. In the HEME program macrocytic anemia is one of the 13 subgroups in which the group of hematologic diseases may be partitioned. The group consists of the 13 different diseases listed in Table 4. Possibly 25 tests deserve consideration to gather information about what member of the macrocytic anemia group exists in the patient.

In determining the optimal nonzero order and zero order rule, only the last ten tests listed in Table 5 are of

interest. The costs associated with each of the first 15 tests can be made equal to 0, and as a consequence there is no problem of choice concerning the performance of these tests in order to collect further information about the patient's physical condition. In practical applications the updated probability distribution of the diseases after having observed the outcomes of the first 15 tests may therefore be taken as the a priori probability distribution of the diseases to start the analysis. In the example that we have worked out, the nonupdated a priori probability distribution as specified in HEME is taken as the starting point.

In estimating the costs of performing each of the last 10 tests in Table 5, the following main cost components were distinguished and estimated separately.

1. The (technical) costs of performance of the test,
2. The costs of possible morbidity and mortality to the patient,
3. The costs of time in the hospital,
4. The costs of days out of the patient's life.

The total costs of performing the tests were obtained by adding up the cost of its main components. The costs of 1) can be considered as mainly independent of the patient involved. In estimating these costs normal charges for laboratory fees, etc., were taken.

The costs of 2), 3), and 4) obviously depend on the particular patient. The age, sex, and physical condition of the patient generally determine the possible morbidity and mortality of the test to the patient and the expected time of his stay in the hospital as a consequence of performing the test. Both the costs charged by the hospital and the costs of days out of the patient's life are determined by the length of this time. In estimating the costs of 2), 3), and 4), the average patient has been taken as a starting point. In most cases a further partition of the main components into subcomponents in determining the costs was applied.

In specifying the costs of misdiagnosing disease j when the patient has disease i , $i \neq j$, for each pair of diseases i and j , the cost consequences of the answers to the following questions have been estimated:

1. Is there a treatment available that can cure or substantially improve the condition of the patient if the correct diagnosis is made?
2. Does the treatment given now, as a consequence of the wrong diagnosis, do serious damage to the patient?
3. If the right treatment is delayed or the harmful treatment is administered, is the effect reversible?

The estimations of the costs of performing the tests and those resulting from misdiagnosis, in dollar units, are specified in Table 5 and Table 6, respectively. The relevant conditional probabilities were derived from

Table 5 Relevant tests in determining the optimal strategy.

No.	Test	Number of Outcomes	Costs
1	x-ray small intestine	4	600
2	biopsy intestine blunted villi	2	900
3	response to B ₁₂ or liver extract	2	2000
4	response to folic acid	2	2000
5	response to pyridoxine	2	2000
6	Schilling test	4	950
7	serum acid phosphatase	2	50
8	serum vitamin B ₁₂	3	50
9	achlohydria after histamine	2	200
10	malabsorption (by serum carotene or xylose test)	2	50

Table 6 Decision losses.

Disease diagnosed	Disease existing in the patient												
	1	2	3	4	5	6	7	8	9	10	11	12	13
1	0	0	10000	10000	1000	10000	10000	10000	10000	50000	10000	1000	0
2	0	0	10000	1000	10000	10000	10000	10000	50000	10000	1000	0	0
3	0	0	0	10000	1000	10000	10000	10000	10000	50000	10000	1000	0
4	2000	2000	10000	0	10000	10000	10000	10000	10000	50000	10000	10000	2000
5	10000	10000	10000	10000	0	1000	0	1000	1000	50000	1000	1000	10000
6	50000	50000	50000	50000	1000	0	0	1000	1000	50000	1000	1000	50000
7	100000	100000	100000	10000	1000	1000	0	1000	1000	50000	1000	1000	100000
8	50000	50000	50000	50000	1000	1000	0	0	1000	50000	1000	1000	50000
9	50000	50000	50000	50000	1000	1000	0	10000	0	50000	1000	1000	50000
10	50000	50000	50000	50000	1000	1000	1000	10000	10000	0	1000	1000	50000
11	25000	25000	25000	25000	1000	1000	0	1000	1000	50000	0	1000	25000
12	100000	100000	100000	100000	10000	10000	10000	10000	10000	50000	10000	0	100000
13	0	0	10000	10000	1000	10000	10000	10000	10000	50000	10000	1000	0

HEME and are shown in Table 7. Finally, Table 8 shows the a priori probability distribution of the diseases that was taken as the starting point in our analysis.

In Fig. 2, a part of the optimal nonzero order rule reached has been laid out. The optimal zero order strategies for groups (sets of tests of fixed length) are given in Table 9 for lengths 1 through 10. As can be derived from Table 9, the overall optimal zero order rule is to perform the tests T₇, T₈, and T₁₀ in an arbitrary sequence.

As mentioned before, the computer program to determine the optimal nonzero order and zero order decision rules has been written in APL. To find the optimal decision rules the program was run for about 160 minutes of CPU time. However, since APL is an interpretive language, the program, rewritten in another language, e.g., PL/I, may be expected to run about 15 or 20 times faster.

Discussion of assumptions

Previously we gave a survey of the underlying assumptions in our analysis. We now consider briefly the consequences of dropping or easing these assumptions.

Assumptions 1 and 2 The first two assumptions are normally met in real life situations and therefore need no further comment.

Assumption 3 If assumption 3 is not satisfied, only a restricted number of all possible test sequences may be feasible. Let, for example, some test T_{*} be available only during the time interval [τ_i, τ_{i+1}], and let us assume that the performance of the first test starts at t = 0 and that each test of the sequence is performed immediately after the outcome of the preceding one has been observed. Then T_{*} is only allowed to be performed as the rth test of a nonzero order decision rule if

$$\tau_i \leq \sum_{j=1}^{r-1} t_j \leq \tau_{i+1} - t_*$$

holds, where t_{*} stands for the performance time of T_{*}. Test T_{*} is only allowed to be performed as the rth test in a zero order decision rule if

$$\tau_i \leq \sum_{j=1}^r t'_j \leq \tau_{i+1} - t'_*$$




Figure 2 Tree for diagnostic problem.

On the one hand, computations become more complicated because the execution times of the tests have to be taken into consideration; on the other hand, a number of possible sequences of tests need not be considered at all.

Assumption 4 If assumption 4 does not hold, then the various diseases do not necessarily form a set of mutually

exclusive events. As a consequence Bayes' formula, to compute the updated probabilities of the diseases after some set of findings has been observed, has to be corrected as was shown earlier [see Eq. (2)]. To restrict computation time the higher order product terms in Eq. (2) may be neglected. Further analysis remains unaffected.

Table 7 Conditional probabilities.

Test attributes	Diseases												
	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	d_9	d_{10}	d_{11}	d_{12}	d_{13}
T_1	$e_{1,1}$	0.95	0.95	0.49	0.19	0.95	0.75	0.95	0.95	0.95	0.98	0.95	0.98
	$e_{1,2}$	0.01	0.01	0.01	0.01	0.01	0.20	0.01	0.01	0.01	0.01	0.01	0.01
	$e_{1,3}$	0.04	0.04	0.50	0.05	0.04	0.05	0.04	0.04	0.04	0.01	0.04	0.005
	$e_{1,4}$	0.00	0.00	0.00	0.75	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.005
T_2	$e_{2,1}$	0.98	0.98	0.98	0.98	0.98	0.50	0.98	0.98	0.98	0.98	0.98	0.98
	$e_{2,2}$	0.02	0.02	0.02	0.02	0.02	0.50	0.02	0.02	0.02	0.02	0.02	0.02
T_3	$e_{3,1}$	0.00	0.00	0.00	0.00	0.80	0.80	0.80	0.80	0.80	0.80	1.00	0.00
	$e_{3,2}$	1.00	1.00	1.00	1.00	0.20	0.20	0.20	0.20	0.20	0.20	0.00	1.00
T_4	$e_{4,1}$	0.30	0.30	0.30	0.30	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.30
	$e_{4,2}$	0.70	0.70	0.70	0.70	1.00	1.00	1.00	1.00	1.00	1.00	0.00	0.70
T_5	$e_{5,1}$	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.90	1.00
	$e_{5,2}$	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.00
T_6	$e_{6,1}$	0.01	0.01	0.01	0.01	0.98	0.98	0.98	0.98	0.98	0.98	0.98	0.05
	$e_{6,2}$	0.99	0.99	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.70
	$e_{6,3}$	0.00	0.00	0.01	0.70	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.20
	$e_{6,4}$	0.00	0.00	0.97	0.28	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.05
T_7	$e_{7,1}$	0.99	0.99	0.99	0.99	0.10	0.10	0.10	0.10	0.10	0.10	0.70	0.99
	$e_{7,2}$	0.01	0.01	0.01	0.01	0.90	0.90	0.90	0.90	0.90	0.90	0.30	0.01
T_8	$e_{8,1}$	0.99	0.99	0.99	0.99	0.02	0.02	0.02	0.02	0.02	0.02	0.30	0.99
	$e_{8,2}$	0.01	0.01	0.01	0.01	0.68	0.68	0.68	0.68	0.68	0.68	0.60	0.01
	$e_{8,3}$	0.001	0.001	0.001	0.001	0.30	0.30	0.30	0.30	0.30	0.30	0.10	0.001
T_9	$e_{9,1}$	0.001	0.001	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.50
	$e_{9,2}$	0.999	0.999	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.50
T_{10}	$e_{10,1}$	0.05	0.05	0.10	0.05	0.05	0.99	0.05	0.05	0.05	0.05	0.05	0.05
	$e_{10,2}$	0.95	0.95	0.90	0.95	0.95	0.01	0.95	0.95	0.95	0.95	0.95	0.95

Assumption 5 If assumption 5 is not satisfied, two or more diseases have to be taken together as an artificial disease, and further analysis remains unaffected.

Assumptions 6, 7, 8, and 9 According to assumptions 6, 7, 8, and 9, the decision model is supplied with all the necessary data to reach a rational decision. Without all these data, the algorithms described have no practical meaning.

Assumption 10 If assumption 10 is not satisfied, uncertainties about test results or unreliable tests have to be accommodated by the underlying model.

Assumption 11 If assumption 11 is not met, properties 1 and 2 generally will not hold. This means also that in searching for the optimal zero order decision rule, distinction has to be made between sequences of the same tests ordered in a different way. Dropping assumption 11, therefore, implies increasing computation time to determine the optimal zero order rule. Note that the number of computations that have to be carried out to determine the optimal nonzero order decision rule does not increase.

Assumption 12 If assumption 12 is not met, then there exists no straightforward way to compare the costs of further testing with those of ceasing further testing and accepting the best diagnosis on the basis of the information collected so far. The decision problem will then have to be formulated in another way; e.g., search for the decision rule (zero order or nonzero order) for which total expected costs of misdiagnosis are minimal and which satisfies the constraint that the expected total time of performing tests is less than or equal to some specified upper bound.

Assumption 13 In practice a physician often starts applying treatments to a patient long before he has taken a final decision with respect to what disease the patient has. The extent to which these treatments can also be considered to be tests makes no difference for our analysis. However, the patient's condition is likely to change as a consequence of each treatment, and it may be necessary to specify the conditional probability of observing some test outcome e_{jq} conditional upon the treatments (tests) already performed on the patient. This implies dropping assumption 11, the consequences of which have already been discussed.

Table 8 Disease a priori probabilities.

Disease	A priori probability
d_1	0.44
d_2	0.01
d_3	0.03
d_4	0.01
d_5	0.01
d_6	0.35
d_7	0.05
d_8	0.015
d_9	0.01
d_{10}	0.01
d_{11}	0.005
d_{12}	0.05
d_{13}	0.01

Table 9 Optimal zero order rules for test sequences of fixed length.

Length	Optimal rule	Expected costs
1	T_8	1099.400
2	T_8, T_{10}	892.845
3	T_7, T_8, T_{10}	752.854
4	T_7, T_8, T_9, T_{10}	899.724
5	$T_1, T_7, T_8, T_9, T_{10}$	1417.450
6	$T_1, T_6, T_7, T_8, T_9, T_{10}$	2308.053
7	$T_1, T_2, T_6, T_7, T_8, T_9, T_{10}$	3206.704
8	$T_1, T_2, T_4, T_6, T_7, T_8, T_9, T_{10}$	5168.769
9	$T_1, T_2, T_3, T_4, T_6, T_7, T_8, T_9, T_{10}$	7148.128
10	$T_1, T_2, T_3, T_4, T_5, T_6, T_7, T_8, T_9, T_{10}$	9148.128

Summary

In this paper the search for a solution to the medical diagnostic problem has been interpreted as a decision rule in stochastic programming. Nonzero order decision rules and zero order decision rules can be seen as pure forms of two different search strategies. Also mixtures of both strategies can well be imagined. The solution to the problem is then: Perform, in an arbitrary order, a sequence of tests of length r , observe the r test outcomes, and decide, conditional upon this information, either to proceed with testing with some next sequence of tests, once again in an arbitrary order, or to stop further testing to accept the best terminal decision, and so on. In a pure nonzero order rule each test sequence has a length 1. Algorithms, quite similar to those discussed for finding the optimal pure nonzero order rule and pure zero order rule, may be applied to solve the mixed strategy case.

Acknowledgments

For helpful suggestions and interesting discussions on the subject of this paper, the author is greatly indebted to

B. J. Flehinger, IBM Thomas J. Watson Research Center, Yorktown Heights, NY. Moreover he thanks Dr. R. L. Engle, Jr., Cornell University Medical College and the New York Hospital, New York, NY, for his willingness in supplying estimations of the costs of misdiagnoses and the costs of performing the tests with respect to the test problem.

References

1. A. Charnes and W. W. Cooper, "Deterministic Equivalents for Optimizing and Satisfying under Chance Constraints," *Oper. Res.* **11**, No. 1, 18 (1963).
2. *Applied Statistical Decision Theory*, H. Raiffa and R. Schlaifer, eds., MIT Press, Boston, 1961.
3. G. H. Symonds, "Chance-Constrained Equivalents of Some Stochastic Programming Problems," *Oper. Res.* **16**, No. 6, 1152 (1968).
4. S. J. Garstka and R. J. B. Wets, "On Decision Rules in Stochastic Programming," *Math. Program.* **7**, 117 (1974).
5. *Probabilistic Programming*, S. Vajda, ed., Academic Press Inc., New York, 1972.
6. A. Charnes and W. W. Cooper, "Chance-Constrained Programming," *Manage. Sci.* **6**, 73 (1969).
7. M. J. Eisner, R. S. Kaplan, and J. V. Soden, "Admissible Decision Rules for the E-Model of Chance-Constrained Programming," *Manage. Sci.* **17**, 337 (1971).
8. B. J. Flehinger and R. L. Engle, Jr., "HEME: A Self-Improving Computer Program for Diagnosis-Oriented Analysis of Hematologic Diseases," *IBM J. Res. Develop.* **19**, 557 (1975).
9. R. L. Engle, Jr., B. J. Flehinger, S. Allen, R. Friedman, M. Lipkin, B. J. Davis, and L. L. Leveridge, "HEME: A Computer Aid to Diagnosis of Hematologic Disease," *Bull. New York Acad. Med.* **52**, No. 5, 584 (1976).
10. G. A. Gorry and G. O. Barnett, "Experience with a Model of Sequential Diagnosis," *Comput. Biomed. Res.* **1**, 490 (1968).
11. *A System for Computer-aided Diagnosis*, G. A. Gorry, ed., Project MAC, Massachusetts Institute of Technology, Cambridge, MA, 1967.
12. *Measuring the Effectiveness of Medical Decisions*, S. Barnoon and H. Wolve, eds., Charles C. Thomas Publishing Co., Springfield, IL, 1972.
13. B. J. McNeil, D. E. Keeler, and S. J. Adelstein, "Primer on Certain Elements of Medical Decision Making," *N. E. J. Med.* **211** (July 31, 1975).
14. W. B. Schwartz, G. A. Gorry, J. P. Kassirer, and A. Essig, "Decision Analysis and Clinical Judgement," *Amer. J. Med.* **55**, 459 (1973).

Received July 16, 1976; revised April 6, 1977

The author, who is on the Faculty of Economics at the University of Groningen, The Netherlands, did the work reported in this paper while a Visiting Professor at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York.