

Lead: A Methodology for Learning Efficient Approaches to Medical Diagnosis

Q1

Saif J. Fakh and Tapas K. Das

Abstract—Determining the most efficient use of diagnostic tests is one of the complex issues facing medical practitioners. With the soaring cost of healthcare, particularly in the US, there is a critical need for cutting costs of diagnostic tests, while achieving a higher level of diagnostic accuracy. This paper develops a learning based methodology that, based on patient information, recommends test(s) that optimize a suitable measure of diagnostic performance. A comprehensive performance measure is developed that accounts for the costs of testing, morbidity, and mortality associated with the tests, and time taken to reach diagnosis. The performance measure also accounts for the diagnostic ability of the tests. The methodology combines tools from the fields of data mining (rough set theory, in particular), utility theory, Markov decision processes (MDP), and reinforcement learning (RL). The rough set theory is used in extracting diagnostic information in the form of rules from the medical databases. Utility theory is used in bringing various nonhomogenous performance measures into one cost based measure. An MDP model together with an RL algorithm facilitates obtaining efficient testing strategies. The methodology is implemented on a sample problem of diagnosing solitary pulmonary nodule (SPN). The results obtained are compared with those from four alternative testing strategies. Our methodology holds significant promise to improve the process of medical diagnosis.

Index Terms—Markov decision processes (MDP), medical diagnosis, reinforcement learning, rough sets.

I. INTRODUCTION

DIAGNOSTIC tests are performed in order to extend the information base beyond what is usually obtained from physical examination and medical history of a patient. The added information allows the physician to obtain a deeper insight of the patient's medical condition. New and advanced diagnostic tests and tools are constantly being introduced to facilitate better understanding and treatment of disease, injury, and congenital or acquired abnormalities.

Diagnosis of a disease and its treatment are not separate processes. Although proper diagnosis helps to narrow the appropriate treatment choices, often treatment is pursued before a diagnosis is reached. This is because the diagnostic process often takes time, and could make the patient's outcomeworse if left unchecked. The choice of diagnostic tests is also affected by the cost and the effect of the testing procedure on the patient. Thus, in the course of patient management, one needs to carefully evaluate the benefits of possible diagnostic steps with

Manuscript received May 11, 2004; revised March 30, 2004, November 5, 2004, December 14, 2004. This work was supported in part by the National Science Foundation Grant # ITR: DMI0113946.

The authors are with the Department of Industrial and Management Systems Engineering, University of South Florida, Tampa, FL 33620 USA.
Digital Object Identifier 10.1109/TITB.2005.855538

Q2

regards to the overarching theme of patient's well being. This can be achieved by optimizing a comprehensive performance measure defined as a function of the costs of the tests, the risk and discomfort associated with the tests, the diagnostic ability of the tests, and the time taken by the process.

A diagnostic process has three distinct phases.

- 1) Based on the initial symptoms and examinations, the physician postulates one or more diagnostic hypotheses and chooses to pursue a subset of these.
- 2) The results from various tests are used to confirm the selected hypothesis. Depending on the circumstance, there could be a wide variety of tests available for the physician to choose from.
- 3) Using the results of the tests, the physician either accepts or discards the selected hypothesis.

Our main focus is on the second phase, where we consider the scenario in which the physician, based on the selected hypothesis, adopts a testing strategy for its confirmation. This paper presents a novel methodology (LEAD) that can be used to derive efficient testing strategies.

Q3

II. BRIEF LITERATURE REVIEW

Medical informatics has become an integral part of successful medical institutions [1], [2]. It uses the technologies developed in the new interdisciplinary field of knowledge discovery in databases (KDD) or data mining, which encompasses tools from areas such as statistics, pattern recognition, machine learning, and data visualization. In the early eighties, the main focus of medical informatics was supporting diagnostic decision making in specialized medical domains through expert systems. One such pioneering system is MYCIN [3], which was followed by numerous other efforts [4], [5].

Pawlak [6], [7], introduced *rough set theory* in the early 1980s as a tool for representing and reasoning with imprecise or uncertain information. Since its introduction, various applications of rough set theory have been developed in the field of medical informatics. Examples include data analysis of preoperative information of patients with duodenal ulcer treated with highly selective vagotomy [8], [9], induction of prognostic rules by analyzing a database on women with breast cancer to determine short-term and long-term follow-up survival [10], development of a prototype expert system for assessing preterm birth risk [11], and diagnosis of progressive encephalopathy in children [12], and acute appendicitis [13], [14]. The literature also documents methods that were developed to optimize the diagnostic process. Some examples of such methods are the analytical hierarchy process (AHP), a multicriteria decision making method that assists physicians in sequentially

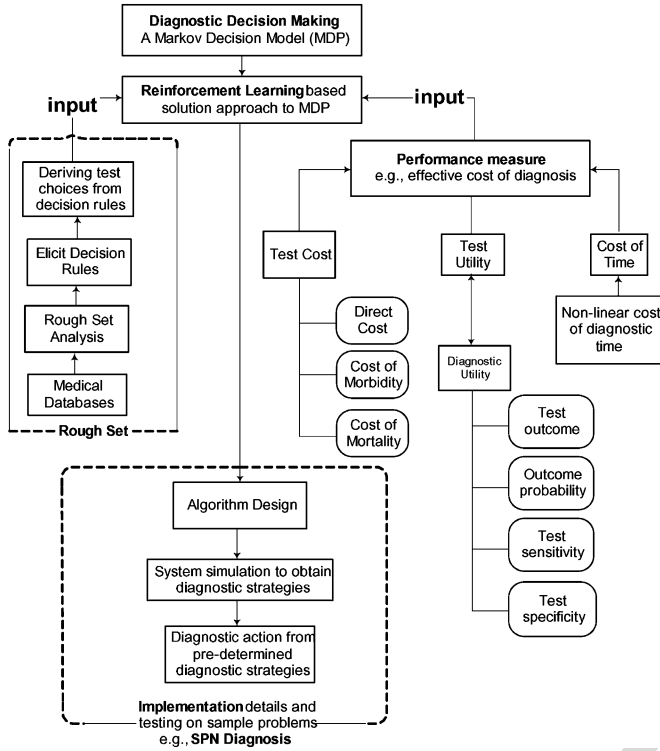


Fig. 1. Overview of LEAD framework.

selecting the most appropriate tests for medical diagnosis [15]; use of modified Bayes' formula to analyze the sixteen test combinations for the diagnosis of Hepatolenticular Degeneration or Wilson's disease [16]; and application of rough sets and Boolean reasoning to select test and procedures for diagnosis of coronary artery disease [17]. Evolutionary algorithms designed to improve the efficiency of healthcare institutions can be found in applications such as DIAPRO-diagnostic process optimization framework [18].

III. LEAD FRAMEWORK

The LEAD framework (shown in Fig. 1) has four major modules that accomplish the tasks needed to arrive at efficient testing strategies. The rough sets module extracts decision rules from medical databases. These rules are then matched with the patient state in order to determine possible test choices (actions). The performance measure module calculates a comprehensive cost measure for evaluating the test choices using utility theory. The test choices and their performance measures serve as primary inputs to the Markov decision process (MDP) based decision model for obtaining diagnostic strategies. Several papers exist in the literature that exploit MDP to model problems related to medial diagnosis [19]–[21]. We solve the MDP model using a simulation based machine learning approach known as reinforcement learning (RL).

The RL algorithm developed here is founded on the value iteration algorithm of MDP. Some of the excellent references on RL are [22], [23]. In what follows, the MDP model is presented

first. Thereafter, a detailed account of the theory/algorithm driving each of the modules of the LEAD framework is given.

A. A Markov Decision Process (MDP) Model

Let the number of possible tests available for diagnosis be N . Define the test set as $T = \{1, 2, \dots, N\}$. The random variable denoting the state of the patient at any decision epoch n can be given as a vector

$$X_n \zeta = (\zeta_1, \zeta_2, \dots, \zeta_n; \gamma_1, \gamma_2, \dots, \gamma_n; d),$$

where $\zeta_1, \zeta_2, \dots, \zeta_n$ denotes the nonoverlapping subsets of tests conducted at the decision epochs such that $\zeta_i \subset T$ for all $i = 1, 2, \dots, n$. The elements $\gamma_1, \gamma_2, \dots, \gamma_n$ denote the test results corresponding to $\zeta_1, \zeta_2, \dots, \zeta_n$, respectively. The last element d of the patient state vector contains the diagnostic outcome, which could be a binary or a multivalued variable. The change in the patient state occurs when the results from the most recent action are obtained, and new action is chosen unless the diagnosis is reached. It is assumed that the time between any two decision epochs is exponentially distributed with the parameter dependent on the system state at the first of those two epochs.

Define $\mathbf{X} = \{X_n : n \in \mathbf{N}\}$ as the patient state process, where \mathbf{N} is the set of natural numbers. It can be easily shown that the Markov property given as $P[X_{n+1} = j | X_0, \dots, X_n] = P[X_{n+1} = j | X_n]$ holds. Hence the stochastic process \mathbf{X} is a Markov process. Define a decision process ζ as $\zeta = \{\zeta_n : n \in \mathbf{N}\}$. Then the combined process (\mathbf{X}, ζ) is a Markov decision process.

B. Rough Sets Approach to Deriving Actions From Decision Rules

We first discuss in brief the rough set approach [7] to obtain decision rules from an existing patient data set. We then discuss how decision rules are matched with a patient state to extract action choices. The matching process also helps in identifying the diagnostic outcome at the conclusion of the diagnosis process, which occurs when any one of the decision rules matches with the patient state.

An information system is a pair $\mathcal{A} = (U, A)$, where U is a finite set of objects and A is a finite set of attributes. For each attribute $a \in A$, there exists a mapping $f_a : U \rightarrow V_a$; i.e., V_a is the set of values for U corresponding to attribute a . For example, the set U may contain the medical history of patients (objects) and the attribute set A could consist of all possible tests that could be conducted on the patients. A decision system is an information system where the attribute set is augmented by a diagnostic outcome attribute d , and can be given as $\mathcal{A} = (U, A \cup d)$, where $d \notin A$. The diagnostic outcome attribute can be interpreted as a classifier on the object set U and is provided by a decision maker or an expert such as a physician.

Let $V = \{V_a : a \in A\}$ and $B \subseteq A$ where B is a reduct obtained using rough set analysis. Atomic formulae over the sets B and V are expressed in the form $a = v$ where $a \in B$ and $v \in V_a$. The set $F(B, V)$ denote the least set containing all the atomic formulae over B and V closed with respect to the propositional connectives \wedge (conjunction), \vee (disjunction), and

TABLE I
 SAMPLE DECISION RULES AND TEST OUTCOMES

Test Outcomes				Sample Decision Rules	
Test Type	α	β_1	β_2	$\alpha = 1 \ \& \ \beta_1 = 1$	$\rightarrow d = 1$
Outcome1	1	0	4	$\alpha = 2 \ \& \ \beta_2 = 5$	$\rightarrow d = 2$
Outcome2	2	1	5	$\alpha = 1 \ \& \ \beta_1 = 0 \ \& \ \beta_2 = 5$	$\rightarrow d = 2$
				$\alpha = 2 \ \& \ \beta_1 = 0 \ \& \ \beta_2 = 4$	$\rightarrow d = 1$

\neg (negation). The rules obtained using rough sets are expressed in the form $\Psi \Rightarrow d = v_d$, where $\Psi \in F(B, V)$ and $v_d \in V_d$, the set of all possible diagnostic outcome attribute values. The elements of a decision rule are Ψ and $d = v_d$, which are referred to as the predecessor and successor, respectively. The predecessor or the conditional part of the decision rule consists of one or more atomic formulas from the set $F(B, V)$ and are used to define actions for a given patient state.

We can rewrite the patient state at a decision epoch n as $X_n = ((\zeta_1, \gamma_1), (\zeta_2, \gamma_2), \dots, (\zeta_n, \gamma_n); d)$, which can be represented as (F_x, d) , where $F_x \subseteq F(B, V)$. Let $B_x \subseteq B$ be the attribute set present in a patient state x and $B_r \subseteq B$ be the attribute set present in the predecessor of a decision rule $r \in R$, where R is the rule set obtained from rough set analysis. The actions for the given patient state x are defined as follows. The attribute set B_x is individually compared with B_r for each $r \in R$.

For a rule r , let v_{a_x} and v_{a_r} denote the values of an attribute a that belongs in both B_x and B_r , i.e., $a \in B_x \cap B_r$. If $v_{a_x} \neq v_{a_r}$ for any $a \in B_x \cap B_r$ then the rule r is discarded, else an action consisting of attribute set $\{B_r \setminus B_x\}$ is added to the action set for the given patient state. For any patient state, if there exists a rule in the rule set such that $v_{a_x} = v_{a_r}$ for all $a_x \in B_x$ and $a_r \in B_r$, where $B_r \subseteq B_x$, it signifies that the diagnosis is reached.

1) *An Example for Deriving Actions from Decision Rules:* Consider a sample problem consisting of three tests: α , β_1 , and β_2 . Possible test combinations with α as the initial test are α , $\alpha\beta_1$, $\alpha\beta_2$, and $\alpha\beta_1\beta_2$. As shown in Table I, a sample decision rule set can be obtained considering two possible outcomes for each test. The diagnostic outcome variable (d) is considered to have three possible values (0: no decision; 1: disease present; 2: disease absent).

Now, consider a patient state $x = (\alpha; 1; 0)$, for which B_x is α . Clearly all of the decision rules that contain the attribute α of the patient state and should be considered. For the first rule from Table I, B_r would be (α, β_1) . Now we have that $B_x \cap B_r = \alpha$. Since $v_{a_x} = v_{a_r} = 1$, the first rule results in an action choice of $B_r \setminus B_x = \beta_1$. The same procedure can be followed for the other rules. For example, the third rule would result in action choice of $B_r \setminus B_x = (\beta_1, \beta_2)$, whereas the second and the fourth rules will be discarded as $v_{a_x} \neq v_{a_r}$ for the attributes in the set $B_x \cap B_r$. Consider another patient state $(\alpha, \beta_2; 2, 5; 2)$. For the second rule, we have that $B_r = (\alpha, \beta_2) = B_x$, $B_r \subseteq B_x$, and $v_{a_x} = v_{a_r}$ for all attributes $a_r \in B_r$, which indicate that a diagnosis is reached.

C. A Utility Theory Based Assessment of Diagnosis Cost

The cost of taking an action is the combined effective cost U.S. dollars (\$) of all tests included in the action. Effective cost

is defined as the actual amount of money spent per unit of diagnostic performance [24]. It is defined as

$$\text{Effective cost} = \frac{\text{Expected Direct Cost}}{\text{Diagnostic Utility}}.$$

Direct Cost of a Test:

- 1) The direct cost includes cost (\$) of conducting the test and the cost (\$) of dealing with the associated risk of morbidity and mortality. The cost of morbidity resulting from a diagnostic test can be estimated as $\sum C_j P_j$, where C_j is the cost of managing the adverse condition j , and P_j is the probability of that condition resulting from the procedure. The cost of mortality can be obtained by multiplying average payoff value of a typical life insurance policy or the cost of a wrongful suit C_m by the probability P_m of mortality resulting from the test. Thus, the expected direct cost can be expressed as [24]

$$\begin{aligned} \text{Expected Direct Cost} &= \text{cost of the test} \\ &+ \sum C_j P_j + C_m P_m. \end{aligned}$$

Diagnostic Utility of Test:

- 2) Any diagnostic test has four traditional outcomes: true positive (TP), true negative (TN), false positive (FP), and false negative (FN). The utility of an outcome is a measure of how desirable is the outcome to a patient and the physician compared with the other outcomes [24]. The utility values can be normalized, for example, on a range between -1 (the worst possible outcome) and $+1$ (the best outcome). A utility value of zero represents a test outcome to which both patient and doctor are indifferent. The scales like -1 to $+1$ have the advantage of clearly demarking the undesirable and indifferent outcomes from the desirable ones. The diagnostic utility (DU) is obtained as the probability weighted sum of all the outcome utilities of a diagnostic test, as given below in the following.

$$\begin{aligned} \text{DU} &= U(\text{TP})P(\text{TP}) + U(\text{TN})P(\text{TN}) \\ &+ U(\text{FP})P(\text{FP}) + U(\text{FN})P(\text{FN}), \end{aligned}$$

where, $U(\cdot)$ = utility of an outcome and $P(\cdot)$ = probability of an outcome. The probabilities are obtained as follows.

$$\begin{aligned} P(\text{TP}) &= P(D)\text{Se}, \\ P(\text{TN}) &= (1 - P(D))\text{Sp}, \\ P(\text{FP}) &= (1 - P(D))(1 - \text{Sp}), \\ P(\text{FN}) &= P(D)(1 - \text{Se}), \end{aligned}$$

where $P(D)$ denotes the probability of disease being present in the patient undergoing the test, and the sensitivity (Se) and specificity (Sp) are given as

$$\text{Se} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

and

$$\text{Sp} = \frac{\text{TN}}{\text{TN} + \text{FP}}.$$

For normalized utility values of -1 to $+1$, a DU value of 1 indicates that a dollar spent for a test buys a full dollar's worth of diagnostic performance. A decrease in the DU value indicates decreasing worth of a dollar spent, which reaches zero for a completely nondiscriminating test. Any DU value less than zero means an additional loss beyond the dollar spent for the test.

LEAD implementation requires as input, for each test, the information needed to calculate the expected direct cost (EDC) and the utility and the probability values for the four test outcomes. The utility values, assigned by the physician, are dependent on the disease scenario, and hence are identical for all available tests [24]. LEAD computes the DU value and the corresponding effective cost for each test. For tests with DU values of zero or less, LEAD ignores the EDC and assigns a very high effective cost. This procedure automatically makes the action choices containing tests with zero or negative DU unworthy of inclusion in the diagnostic strategy.

D. Simulation Based Learning Scheme

As shown in Fig. 2, the scheme begins with a simulated arrival of a new patient. Based on the information collected from the patient and the past learning experience (in the form of R-values), an action to perform one or more tests is chosen. Outcomes for the tests are simulated by sampling from probability distributions that characterize the tests. If diagnosis is not reached, the new patient state is determined and the next action is chosen. If diagnosis is reached, the R-values for all the actions selected during the diagnosis process of the patient are updated, and the scheme repeats by simulating a new patient. An algorithmic version of the scheme is presented as follows.

Define trajectory as a path along which a patient traverses during the diagnosis process. It begins with a new patient arrival and ends when a diagnosis is reached. Let $S = (x_1, x_2, \dots, x_n)$ denote the set of states visited in a trajectory, where x_n represents the patient state at the n th decision epoch of the trajectory. Let $\tilde{A} = (\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_n)$ denote the set of actions in states x_1, x_2, \dots, x_n of S and $C = (c_1(x_1, \tilde{a}_1), c_2(x_2, \tilde{a}_2), \dots, c_n(x_n, \tilde{a}_n))$ denote the set of immediate costs at the decision epochs within a trajectory. Also, let m, T and $maxruns$ denote the indices for count of decision epochs over all trajectories, the count of total number of trajectories, and the maximum number of trajectories that are to be simulated during the learning process respectively. The patient state space E is initialized with starting states only, and is appended with new states that are visited along the trajectories. For each starting state and the new states appended to E , corresponding R-values are initialized as $R(j, \tilde{a}) = 0$, where $\tilde{a} \in \tilde{A}_j$ and \tilde{A}_j is the set of actions available in state $j \in E$. Initialize sets S, \tilde{A}, C as empty sets. Choose appropriate values of $\alpha_0, \beta_0, \gamma_0, \alpha_\tau, \beta_\tau, \gamma_\tau$, which are initial values and the decay parameters for the learning (α and β) and exploration (γ) rates. Define ρ_i, ρ_f, m_t , and $stepcost$ as temporary variables. Set $\rho_i = 0, \rho_f = 0, n = 1, m = 0, m_t = 0, T = 1, \alpha = \alpha_0, \beta = \beta_0$, and begin the learning phase.

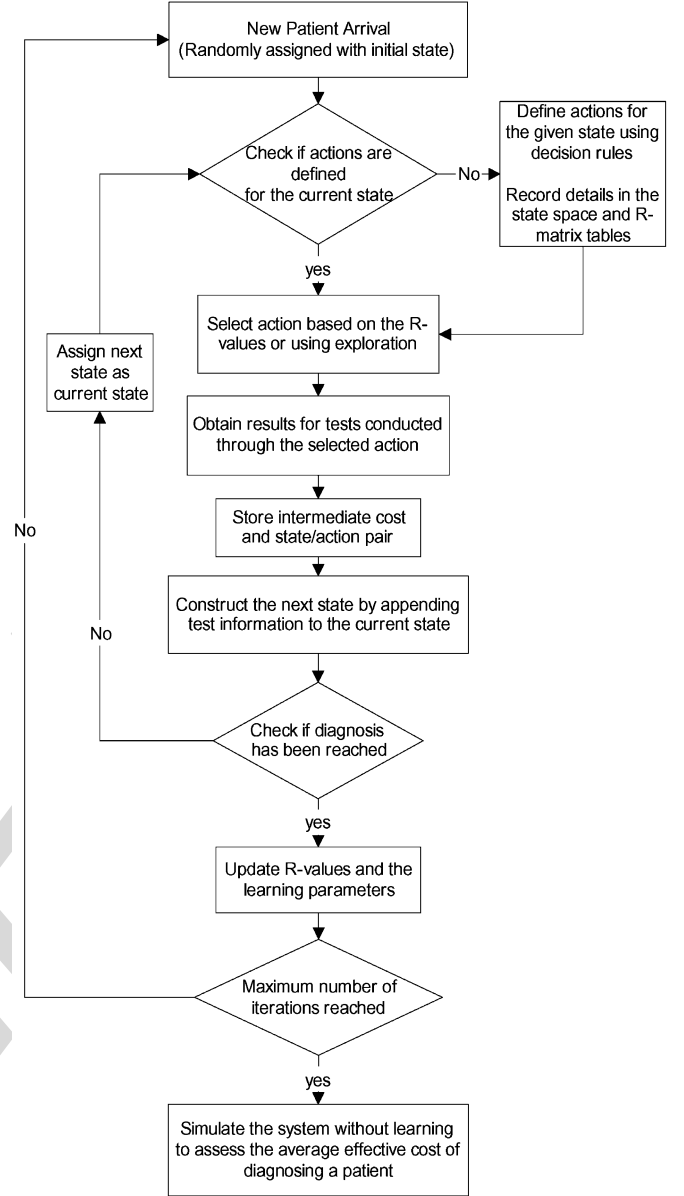


Fig. 2. Detailed flowchart for RL algorithm.

- 1) Simulate a new patient arrival and assign a randomly chosen initial state.
- 2) If the current state has been previously visited, then go to Step 3, else
 - Identify the set of possible actions for the current state by comparing the attributes present in the state vector with the attributes of the individual rules present in the rules set.
 - Append the R array with all combinations of the current state and the possible actions.
- 3) Update γ as

$$\gamma = \frac{\gamma_0}{1 + u}, \quad \text{where } u = \frac{m^2}{\gamma\tau + m}.$$

With probability $(1 - \gamma)$, choose an action $\tilde{a}_n \in \tilde{A}_{x_n}$ for which $R(x_n, \tilde{a}_n)$ is minimum; otherwise choose a random action from the set $\{\tilde{A}_{x_n} \setminus \tilde{a}_n\}$.

- 4) Obtain results for the test(s) included in the action \tilde{a}_n .
- 5) Append the current state (x_n), action chosen (\tilde{a}_n), and the immediate cost $c_n(x_n, \tilde{a}_n)$ as the n th element of the vectors S , \tilde{A} , and C , respectively.
- 6) Construct the next state by appending the test types and their results to the current state vector. Also, check if the next state indicates that a diagnosis is reached.
- 7) If diagnosis has been reached, then
 - Obtain cost per step within the trajectory (denoted as *stepcost*) as

$$\text{stepcost} = \left(\sum_{k=1}^n C_k + c_1 + c_2 * n + c_3 * n^2 \right) / n$$

where the quadratic function above is used to access the cost of time involved in the steps needed to reach diagnosis; c_1 , c_2 , and c_3 are constants that should be chosen suitably.

- Update average effective cost and the R-values of the current state as follows.
For $k = 1$ to n

$$R_{\text{new}}(x_k, \tilde{a}_k) = (1 - \alpha)R_{\text{old}}(x_k, \tilde{a}_k) +$$

$$\alpha \left[\text{stepcost} - \rho_f + \min_{\tilde{a} \in \tilde{A}_{x_k+1}} R_{\text{old}}(x_k + 1, \tilde{a}_k + 1) \right]$$

$$\rho_f = (1 - \beta)\rho_i + \beta \left[\frac{m_t \rho_i + \text{stepcost}}{m_t + 1} \right].$$

Set $m_t \leftarrow m_t + 1$ and $\rho_i \leftarrow \rho_f$

$$\alpha = \frac{\alpha_0}{1+u}, \quad \text{where } u = \frac{m_t^2}{\alpha_\tau + m_t}$$

$$\beta = \frac{\beta_0}{1+u}, \quad \text{where } u = \frac{m_t^2}{\beta_\tau + m_t}.$$

- If $T = \text{maxruns}$
Go to step 8
Else,
- Reinitialize S, \tilde{A}, C as empty sets.
- Set $m \leftarrow m + 1, T \leftarrow T + 1, \tilde{A}$, and $n = 1$
Go to step 1
Else,
- current state \leftarrow next state
- Set $m \leftarrow m + 1$ and $n \leftarrow n + 1$
- Go to step 2.

- 8) Start a fresh simulation using the existing R-values without learning and exploration for *maxruns* of trajectories to estimate the cost of diagnosing a patient averaged over all trajectories. This is referred to as the learned phase of the simulation

IV. TEST PROBLEM AND IMPLEMENTATION OF LEAD

Diagnosis of solitary pulmonary nodule (SPN) was considered as the test problem for implementation of LEAD. An SPN or ‘‘coin lesion’’ has been defined as a single spherical lesion, 1 to 6 cm in size, completely surrounded by lung without associated atelectasis or adenopathy [25]. The cause of SPN can

be a malignant disease such as bronchogenic carcinoma, lymphoma, sarcoma, or it may result from a benign condition such as hamartoma, granuloma, and sclerosing hemangioma. Malignancy is estimated to occur in 20%–40% of the patients with SPN. The diagnostic challenge with SPN is in establishing whether the nodule is benign or malignant. If the SPN is malignant, the physician could either expedite the removal of the nodule or avoid thoracotomy in patients with surgically incurable lung neoplasms. If the SPN is benign, the physician should avoid thoracotomy because benign nodules rarely pose a risk and need not be resected.

The large amount of information and the risks associated with the diagnostic procedures complicates this process. For example, as many as 100 attributes obtained from various diagnostic procedures may be available to the physician to make the decision. Even for experienced physicians, it is usually difficult to effectively synthesize the array of available information. Examples of risk involved with the diagnostic procedures include occurrence of pneumothorax in 15% to 20% of the patients undergoing fine needle aspiration (FNA) procedure, and mortality of 2% to 3% due to thoracotomy. Hence, it is important to consider morbidity and mortality associated with the invasive procedures. However, dependence only on noninvasive procedures could be misleading. For example, 20%–40% of resected lesions are generally found to be benign even after sophisticated computer tomography (CT) imaging.

LEAD is implemented in the diagnosis of solitary pulmonary nodule using the information set (data) containing 50 objects presented in [26]. The information set is not reproduced here as it is not essential to refer to it in order to follow the implementation of LEAD. The methodology requires following information for implementation.

A. Selection of Attributes

In the model formulation, the set $T = \{1, 2, \dots, N\}$ was defined as the set of all tests associated with the diagnosis. Each test can provide one or more attributes that can be utilized in the process of diagnosis. For LEAD implementation on SPN, 18 different attributes as provided in [26] were used, which are presented in Section IV-D.

B. Decision Rules to Define the Action Sets

The rule set utilized in this study was obtained from [26] and is shown in Table II. These rules were obtained through a rough set analysis of the SPN information set that generated the minimal subset of attributes or reducts. The decision attribute values used here are B (benign) and M (malignant).

C. Generation of Test Attribute Values

The values for the attributes that constitute an action were generated from the probability distributions defined using the SPN data obtained from [26, Table XIV]. The data set contains 50 cases of confirmed diagnosis with values for all the 18 attributes in each case. Probability distributions for attributes with discrete values were obtained directly based on the

TABLE II
DECISION RULE OBTAINED FROM ROUGH SET ANALYSIS

Rule #	Decision Rule
1	IF (F10<=1) AND (F17<=89) THEN (d=B)
2	IF (F2 IN [0.9,1.55]) AND (F5<=0) AND (F6>=2) AND (F8 IN [2.75,8.25]) THEN (d=B)
3	IF (F14 >= 78) AND (F17<=54) AND (F18 <= 4.1) THEN (d=B)
4	IF (F14 IN [94,109]) AND (F16 >=103) THEN (d=B)
5	IF (F6<=2) AND (F10>=1) AND (F16<=94) THEN (d=M)
6	IF (F2<=0.9) AND (F10>=1) THEN (d=M)
7	IF (F10>=1) AND (F13<=0) AND (F14<=74) THEN (d=M)
8	IF (F15<=52) THEN (d=M)
9	IF (F14<=80) AND (F17>=31) THEN (d=M)
10	IF (F14 IN [81,87]) THEN (d=M)
11	IF (F6<=2) AND (F10>=1) AND (F14>=84) AND (F15>=63) THEN (d=M)
12	IF (F16<=93) AND (F17>=49) THEN (d=M)
13	IF (F6>=2) AND (F14>=91) AND (F15<=97) THEN (d=M)
14	IF (F2>=1.65) AND (F15>=116) THEN (d=M)
15	IF (F8<=0) AND (F10<=1) THEN (d=B)
16	IF (F8>=0) AND (F9<=0) AND (F18<=4.1) THEN (d=B)
17	IF (F2 IN [0.9,1.1]) AND (F4>=2) AND (F18<=8.25) THEN (d=B)
18	IF (F2 IN [1.1,1.55]) AND (F6>=2) AND (F8<=0) THEN (d=B)
19	IF (F2 IN [2.25,2.75]) AND (F4<=1) THEN (d=B)
20	IF (F2 <=2) AND (F6<=2) AND (F10>=1) THEN (d=M)
21	IF (F2 >=2.75) AND (F18<=8.6) THEN (d=M)
22	IF (F2 <=2.75) AND (F13<=0) AND (F18>=5.3) THEN (d=M)
23	IF (F2 <=3.25) AND (F18>=8.45) THEN (d=M)
24	IF (F2 >=3.75) THEN (d=M)
25	IF (F2 <=6) AND (F4<=1) AND (F6>=2) THEN (d=M)
26	IF (F7 >=1) AND (F8<=0) THEN (d=M)

TABLE III
EFFECTIVE COSTS OF TEST ATTRIBUTES

Test attributes (Attribute identifier)	Effective Cost	
	Set A	Set B
Computed Tomography-radius (F2)	500	650
Computed Tomography-area (F3)	150	150
Computed Tomography-borders (F4)	300	300
Calcification type (F5)	85	850
Location in thorax (F6)	140	260
Nodes (F7)	260	140
Other SUS lesion (F8)	190	190
PET PN image > bg (F10)	450	750
Force expiratory volume, 1 sec, % predicted (F14)	750	450
Adjusted DLCO, % predicted (F15)	600	600
Forced vital capacity, % predicted (F16)	425	425
Forced expiratory flow, 25-75% (F17)	550	800
PET standard update value (F18)	650	500

frequencies for the attribute values in the data set. The same approach was adopted for the continuous valued attributes after they were discretized. The level of discretization should depend on the desired accuracy of the diagnosis process. This issue is discussed further in Section V.

D. Effective Cost for Tests

For the purpose of implementation, the effective cost values for the test attributes were chosen somewhat arbitrarily. Two sets

of such values (denoted as Set A and Set B) are listed in Table III. Some of the attributes that are not included in Table II, such as age (F1), sex (F9), smoking level (F11), BMI (F12), and history of cancer (F13) are obtained from patient's history at no cost.

V. IMPLEMENTATION CONSIDERATIONS

This section discusses the key aspects of implementation of the RL algorithm within the LEAD methodology.

A. Discretization of the Available Dataset

The SPN data set on which LEAD was implemented contains ten continuous attributes and eight discrete attributes. Since the implementation of the MDP based RL algorithm requires a discrete state space, the continuous valued attributes were discretized. The discretized values were used for 1) defining patient states, 2) transforming decision rules, and 3) generating values for various attributes during simulation runs. The number of steps in which the attribute values were discretized was higher for those appearing frequently in the rule set. It may be noted that the number of steps in which to discretize an attribute value range has two significant implications. With finer discretization, the learning time increases due to the increase in patient states, since a higher number of R-values must be learned. However, finer discretization improves the accuracy of the learned diagnosis strategy.

B. Dynamic Generation of the States

The patient states are generated progressively by appending attributes of the tests that are ordered for the patient to an initial attribute vector, which is a 5 tuple. Since tests (attributes) can be chosen in any order, all possible combinations of attribute vectors (of all sizes) must be considered. For example, beyond the five attributes in the initial states, if two new attributes are to be appended, then there are ${}^{13}C_2$ possible ways to do that (since there is a total of eighteen attributes). The patient state also contains the values of the attributes. Hence the number of possible patient states in this example will be $I * {}^{13}C_2 * d_1 * d_2$ where I indicates the number of initial states, and d_1 and d_2 represent the possible values of the new attributes. Let K be the set of all possible attribute vectors of sizes 5 through 18. For any vector $k \in K$ of size n , possible combinations of attribute values can be given as

$$S_k = \prod_{i=1}^n d_i$$

where d_i is the number of distinct values of attribute i and n is the number of attributes present in the vector. The cardinality of the set K is then given by

$$|K| = 1 + \sum_{j=1}^{13} C_j^{13}$$

where 1 indicates the fixed set of initial attributes, and the summation takes care of all the possible ways of appending the

TABLE IV
 VALUES FOR THE LEARNING PARAMETERS

Learning parameter	Initial value	Constant
α	0.1	1000000
β	0.1	250000
γ	0.5	100000

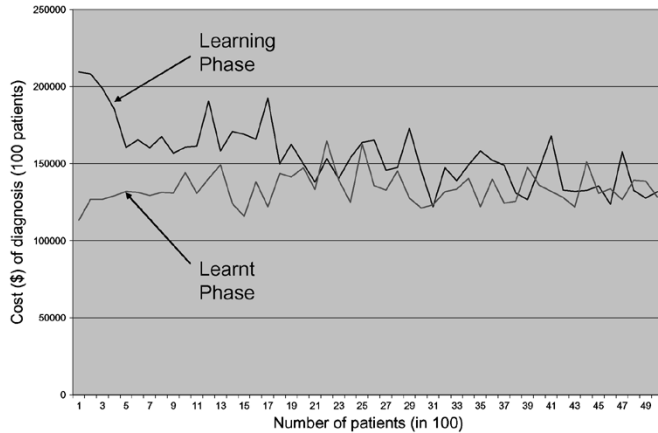


Fig. 3. LEAD results for cost set A.

initial vector. Thus, the state space E is given by

$$|E| = \sum S_k, \quad \text{for all } k \in K.$$

The numerical value of the number of states in the test problem was of the order of 6.69×10^{11} . Hence, dynamic generation of states, as opposed to initializing all 6.69×10^{11} states, reduced a significant amount of computational burden for LEAD.

C. Learning Parameters

Three learning parameters involved in the implementation of the LEAD are α , β , and γ . The running values of the learning parameters at any given decision epoch are determined by their initial values (α_0 , β_0 , and γ_0) and the large constants (α_τ , β_τ , γ_τ) and that are used to control the decay rate. The initial values and the constants are listed in Table IV.

D. Software Framework

The application developed for the implementation of the LEAD methodology is based in an Oracle 8i database. Data management is one of the major challenges in an RL implementation, especially when the state space is large. Oracle provides a resilient O-RDBMS framework for efficient data structuring and storage. The logic for the algorithm provided in Fig. 2 was programmed using PL/SQL.

VI. RESULTS AND ANALYSIS

The numerical results obtained by applying the LEAD solution methodology to the diagnosis of solitary pulmonary nodule are presented here. The results also include a comparative study between the LEAD and several alternative diagnostic strategies. Fig. 3 shows the plot of the learning and the learnt phases of LEAD using cost set A. Similar results from cost set B are

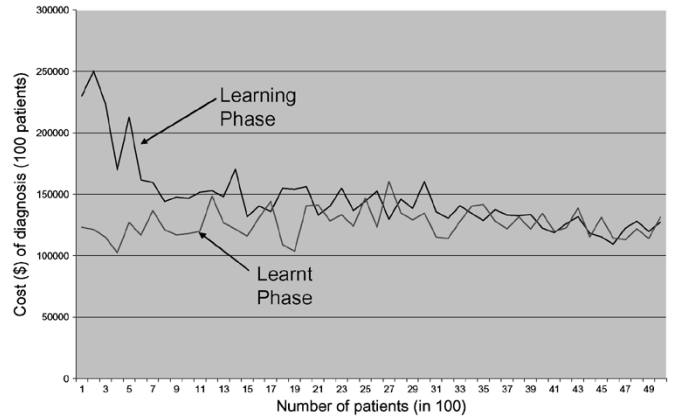


Fig. 4. LEAD results for cost set B.

 TABLE V
 COMPARISON OF COST AVERAGED FOR THE FIRST AND THE LAST 500 PATIENTS SIMULATED DURING LEARNING

	First 500 runs		Last 500 runs		% Reduction in Avg. Cost
	Average Cost (\$)	95% C.I.	Average Cost (\$)	95% C.I.	
Set A	1924.81	90.73	1346.75	98.25	30
Set B	2172	103.60	1213.6	113.93	44

shown in Fig. 4. The plots from the learning phase show the evolution of the diagnostic costs as the learning agent continuously improves its strategy by updating the R-values, whereas the learnt phase plots indicate the cost corresponding to the final set of R-values obtained from the learning phase. The horizontal axis represents the number of patients used in the simulated diagnostic process, and the vertical axis represents the cost of diagnosis. Since a large number of patients (about 5000) were simulated, instead of plotting costs for individual patients, the sum of costs for every 100 consecutive patients is plotted.

The trends of the learning curves clearly indicate a significant policy improvement through a drop in the diagnostic cost as learning proceeds. The policy improvement during the learning process is also evident from the results summarized in Table V, where the cost of diagnosis per patient averaged for the first and the last 500 patients (simulation runs) are shown.

A. Benchmarking LEAD

The policy resulting from LEAD was benchmarked with four other alternative policies. The comparison was based on the total cost of diagnosis per 100 patients. The four policies chosen for benchmarking differ in their methods of selecting actions, which are described as follows.

- 1) Minimum cost policy (HMO policy): The action with minimum effective cost is chosen for any given patient state.
- 2) Random action policy: An action from all available actions is chosen randomly.
- 3) Minimum number of tests policy: The action that involves the minimum number of tests to be performed is selected.

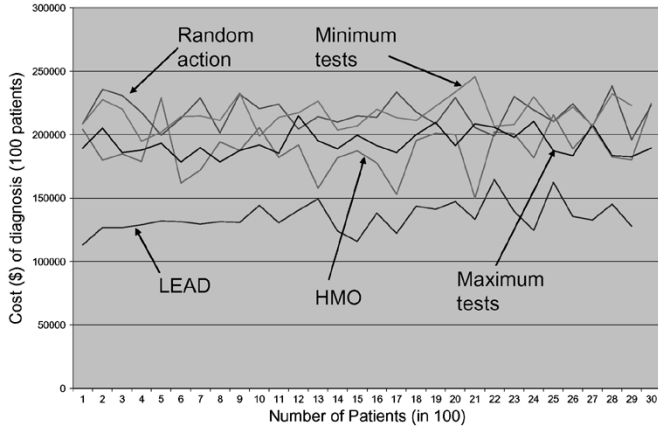


Fig. 5. Diagnostic cost of LEAD and other benchmark policies.

TABLE VI
COMPARISON OF THE TOTAL COST OF DIAGNOSIS PER PATIENT AVERAGED OVER 3000 RUNS

Policy	Average Cost (\$)	95% C.I.	Cost (%) above LEAD
LEAD	1338.52	46.2	N/A
Minimum Cost	1863.37	59.67	28.17
Random Action	2163.61	42.31	38.14
Minimum Tests	2171.91	41.49	38.37
Maximum tests	1932.45	30.99	30.74

TABLE VII
ACTIONS OF HMO AND LEAD POLICIES IN MOST VISITED INITIAL STATES FOR COST SET A

Most visited initial state	HMO		LEAD	
	Action attributes	Action Cost (\$)	Action attributes	Action Cost (\$)
21	7,8	450	7,8	450
20	7,8	450	2	500
5	7,8	450	7,8	450
22	7,8	450	2,18	1150
23	7,8	450	7,8	450

- 4) Maximum number of tests policy: The action that involves the maximum number of tests to be performed is selected.

A visual comparison of the performances for the four alternate policies along with the LEAD policy for cost set A is shown in Fig. 5. As evident from the plots, the results obtained from the LEAD fare better than all the other approaches.

A numerical performance summary of all the five policies is presented in Table VI. It can be seen from the results presented that the benchmark policies cost 28–38% higher than the LEAD policy.

Since the HMO policy performed the best among the alternatives of LEAD, for the purpose of comparison, actions for some of the frequently visited patient states as adapted by HMO and LEAD policies are summarized in Tables VII and VIII. The data in these tables gives a visual depiction of the action choices made by the policies. Actions chosen for five of the most frequently visited initial states during the simulation runs with cost set A are presented in Table VII. The action choices summarized in Table VIII are for some of the most frequently

TABLE VIII
ACTION OF HMO AND LEAD POLICIES IN MOST VISITED INTERMEDIATE STATES FOR COST SET B

Most visited intermediate state	HMO Policy		LEAD policy	
	Action attributes	Action Cost (\$)	Action attributes	Action Cost (\$)
1451	14	450	2,4	950
1711	14	450	2,8	1150
437	7,8	330	17	800
414	7,8	330	6,16	685
869	7,8	330	14	450

visited intermediate states during the simulation runs with cost set B. The first column in each table indicates the patient state identified by a coded number. The actions under the two policies are indicated by their attributes and the corresponding effective costs.

Though for the initial states (Table VII), the LEAD policy often chose the minimum cost action, it may be noted that for intermediate states (Table VIII), the LEAD policy actions are significantly different with much higher immediate effective costs. Given the fact that the LEAD policy provides the lowest average cost of diagnosis, the action choices indicate that the learning agent is able to account for the long term impact of its decisions in the diagnosis process rather than being guided by the immediate cost.

VII. DISCUSSION

A novel methodology for developing efficient strategies for medical diagnosis is presented. The problem of diagnosis is first modeled as a Markov decision process (MDP). The patient state is defined as a vector of attributes of the diagnostic tests conducted and their numerical outcomes. For each patient state, possible actions are obtained from the set of decision rules that are generated by mining the available data set using rough set theory. Since there are a large number of patient states and a significant number of corresponding action choices, a machine learning algorithm is developed for obtaining efficient diagnostic strategies. The resulting methodology (named LEAD) is implemented on a problem of diagnosing solitary pulmonary nodule (SPN), for which data and the corresponding rule set are available in the literature. The results show that the learning agent is successful in significantly improving its strategy by obtaining good R-value estimates during numerous simulation runs (Figs. 3 and 4). The improvement is assessed from the drop in the average cost of diagnosing a patient. Since the learned phase implementation uses the final R-values, it does not encounter the poor strategies of the initial part of the learning phase, and thus obtains a better (lower) estimate of the average diagnosis cost.

As can be seen from Step 7 of the RL algorithm, the learning agent updates the R-values of all its actions for a patient after a diagnosis is reached. (This is known in RL literature as delayed updating.) As a result, the cumulative impact of all the action choices affects updating of R-values enabling the RL agent to assess and incorporate the true impact of its actions. This is

evident from the action choices shown in Tables VII and VIII, and the plots of average cost in Fig. 5. Note that, though the RL agent (LEAD policy) often adopts more expensive test choices, it still maintains a much lower average diagnosis cost compared to the other policies.

VIII. CONCLUSION

Most of the work done in the area of medical decision making considers various learning techniques such as neural networks, decision trees, rough sets, and reinforcement learning in isolation of each other. LEAD combines the principles of rough set theory, utility theory and reinforcement learning to provide a robust mechanism for obtaining efficient strategies for medical diagnosis. The methodology obtains all of its parameters from the databases of previous cases. Thus, it is primarily data driven and requires little interaction with domain experts.

As with the implementation of most of real world decision making models, LEAD could be computationally burdened by very large number of patient states. However, the method of dynamic generation of patient states, which allows LEAD to deal with only the previously encountered states and not the whole state space, coupled with a simulation based RL approach, make this methodology computationally viable. An important feature of LEAD is that it is designed for offline use; i.e., the diagnostic strategies for all patient states are predetermined. The practicing physician needs only to specify the patient state and retrieve the corresponding diagnostic action, and hence is not affected by the computational constraints. The methodology should be rerun periodically to get new diagnostic strategies using the updated databases.

REFERENCES

[1] E. Coiera, *Guide to Medical Informatics, the Internet and Telemedicine*, 1st ed. London, U.K.: Oxford Univ. Press, 1997.

[2] E. H. Shortliffe and L. M. Fagan, *Medical Informatics: Computer Applications in Health Care and Biomedicine*, 2nd ed. New York, NY: Springer-Verlag, 2000.

[3] E. H. Shortliffe, *Computer-Based Medical Consultations: MYCIN*. Amsterdam, The Netherlands: Elsevier, 1976.

[4] R. A. Miller, H. E. Pople, and J. D. Myers, "Internist-I, an experimental computer-based diagnostic consultant for general internal medicine," *New England J. Med.*, vol. 307, no. 8, pp. 468–476, 1982.

[5] H. E. Pople, "Heuristic methods for imposing structure on ill structured problems: The structuring of medical diagnosis," *Artificial Intelligence in Medicine, AAAS Selected Symp. Series*, pp. 119–185, West View Press, 1982.

[6] Z. Pawlak, "Rough sets: Theoretical aspects of reasoning about data," *System Theory, Knowledge Engineering and Problem Solving*, Kluwer, vol. 9, D 1991.

[7] Z. Pawlak, "Rough sets," *Int. J. Inf. Comput. Sci.*, vol. 11, no. 5, pp. 341–356, 1982.

[8] J. Fibak, Z. Pawlak, K. Slowinski, and R. Slowinski, "Rough sets based decision algorithm for treatment of duodenal ulcer by HSV," *Bull. Polish Acad. of Sci., Biological Science*, vol. 34, pp. 227–249, 1986.

[9] K. Slowinski, "Rough classification of HSV patients," in *Intelligent Decision Support: Handbook of Applications and Advances in Rough Sets Theory, System Theory, Knowledge Engineering and Problem Solving*, vol. 11, Norwell, MA: Kluwer, 1992, pp. 77–93.

[10] H. E. Pople, "Heuristic methods for imposing structure on ill structured problems: The structuring of medical diagnosis," in *Artificial Intelligence in Med., AAAS Selected Symp. Series*: West View Press, 1982, pp. 119–185.

[11] L. Woolery and J. Grzymala-Busse, "Machine learning for an expert system to predict preterm birth risk," *J. Amer. Med. Informatics Assoc.*, vol. 1, no. 6, pp. 439–446, 1994.

[12] A. Wakulicz-Deja and P. Paszek, "Diagnose progressive encephalopathy applying the rough set theory," *Int. J. Medical Informatics*, vol. 46, no. 2, pp. 119–127, 1997.

[13] U. Carlin, J. Komorowski, and A. Øhrn, "Rough set analysis of medical datasets in a case of patients with suspected acute appendicitis," in *Proc. ECAI'98 Workshop on Intelligent Data Analysis in Medicine and Pharmacology*, 1998, pp. 18–28, (IDAMAP'98).

[14] U. Carlin, J. Komorowski, and A. Øhrn, "Rough set analysis of patients with suspected acute appendicitis," in *Proc. 7th Conf. Information Processing and Management of Uncertainty in Knowledge-Based Systems*, 1998, pp. 1528–1533, (IPMU'98).

[15] F. Castro, P. Caccamo, J. Carter, B. Erickson, W. Johnson, E. Kessler, N. Ritchey, and C. Ruiz, "Sequential test selection in the analysis of abdominal pain," *Medical Decision Making*, vol. 19, pp. 178–183, 1999.

[16] R. Richards, J. Hammitt, and J. Tsevat, "Finding the optimal multiple-test strategy using a method to Logistic Regression," *Medical Decision Making*, vol. 16, pp. 367–375, 1996.

[17] A. Ohm and J. Komorowski, *Analyzing the Prognostic Power to Cardiac Tests Using Rough Sets*, 1997. [Online]. Available: <http://citeseer.ist.psu.edu/165281.html>.

[18] V. Podgorelec and P. Kokal, "Towards more optimal medical diagnosing with evolutionary algorithms," *J. Medical Systems*, vol. 25, no. 3, pp. 195–219, 2001.

[19] M. Hauskrecht and H. Fraser, "Planning treatment of ischemic heart disease with partially observable Markov decision processes," *Artificial Intell. in Med.*, vol. 18, pp. 221–244, 2000.

[20] O. Alagoz, A. Schaefer, L. Maillart, and M. Roberts, "Determining the optimal timing of living-donor liver transplantation using a Markov decision process (MDP) model," *Medical Decision Making*, vol. 22, 2002.

[21] P. Magni, S. Quaglini, M. Marchetti, and G. Barosi, "Deciding when to intervene: A Markov decision process approach," *Int. J. Medical Informatics*, vol. 60, no. 3, pp. 237–253, 2000.

[22] A. Gosavi, *Simulation-based Optimization: Parametric Optimization Techniques and Reinforcement Learning*. Norwell, MA: Kluwer, 2003.

[23] S. Sutton and A. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.

[24] D. Patton and J. Woolfenden, "A Utility-based model for comparing the cost-effectiveness of diagnostic studies," *Investigative Radiology*, vol. 24, pp. 263–271, 1989.

[25] D. Ost, A. Fein, and S. Feinsilver, "The solitary pulmonary nodule," *New England J. Medicine*, vol. 348, no. 25, pp. 2535–2542, 2003.

[26] A. Kusiak, J. Kern, K. Kernstine, and B. Tseng, "Autonomous decision making: A data mining approach," *IEEE Trans. Inf. Technol. Biomed.*, vol. 4, no. 4, 2000.

Q4



Saif J. Fakih received the M.S. degree in industrial engineering in 2004 from the University of South Florida, Tampa, FL, USA.

Currently, he is working as a Senior Application Developer/Project Manager with the Center for 21st Century Teaching Excellence at the University of South Florida. His areas of research include data mining, medical informatics, and integration of technology with teaching.



Tapas K. Das received the Ph.D. degree in industrial engineering in 1989 from Texas A&M University, College Station, TX.

He is a Professor of industrial and Management Systems Engineering at the University of South Florida, Tampa, FL, USA. His current research includes design of reliable de-regulated electric power markets using game and auction theory; process control and monitoring using sensor data analysis via wavelet based multiresolution analysis approaches; and working with K-12 schools in improving science

education. His research these areas is funded by the National Science Foundation. His research interests also include developing methodologies for medical diagnosis and cancer drug response analysis.

Dr. Das is a Member of the Institute of Industrial Engineers (IIE) and the Institute of Operations Research and Management Sciences (INFORMS).

Queries

- Q1: Pls. Mention IEEE membership.
- Q2: Pls. Need author(s) e-mail add.
- Q3: AU: Please spell out “LEAD”.
- Q4: AU: Pub. and conf. location? of refs. 5, 10, and 13.

IEEE
Proof