Annals of Operations Research 119, 15–42, 2003

Coronary Risk Prediction by Logical Analysis of Data *

SORIN ALEXE salexe@rutcor.rutgers.edu
RUTCOR, Rutgers University, Piscataway, NJ 08854, USA

EUGENE BLACKSTONE
Cleveland Clinic Foundation, 9500 Euclid Avenue, Cleveland, OH 44195, USA

PETER L. HAMMER hammer@rutcor.rutgers.edu
RUTCOR, Rutgers University, Piscataway, NJ 08854, USA

HEMANT ISHWARAN, MICHAEL S. LAUER and CLAIRE E. POTHIER SNADER lauerm@ccf.org
Cleveland Clinic Foundation, 9500 Euclid Avenue, Cleveland, OH 44195, USA

Abstract. The objective of this study was to distinguish within a population of patients with known or suspected coronary artery disease groups at high and at low mortality rates. The study was based on Cleveland Clinic Foundation's dataset of 9454 patients, of whom 312 died during an observation period of 9 years. The Logical Analysis of Data method was adapted to handle the disproportioned size of the two groups of patients, and the inseparable character of this dataset – characteristic to many medical problems. As a result of the study, we have identified a high-risk group of patients representing 1/5 of the population, with a mortality rate 4 times higher than the average, and including 3/4 of the patients who died. The low-risk group identified in the study, representing approximately 4/5 of the population, had a mortality rate 3 times lower than the average. A Prognostic Index derived from the LAD model is shown to have a 83.95% correlation with the mortality rate of patients. The classification given by the Prognostic Index was also shown to agree in 3 out of 4 cases with that of the Cox Score, widely used by cardiologists, and to outperform it slightly, but consistently. An example of a highly reliable risk stratification system using both indicators is provided.

Keywords: classification, data mining, Logical Analysis of Data, partially defined Boolean functions, risk indices, risk prediction

AMS subject classification: 62H30, 06E30, 90C09, 90C10, 90C27, 94C10

1. Introduction

The dataset of the Cleveland Clinic Foundation (CCF) concerning patients referred for symptoms-exercise electrocardiography between September 1990 and March 1998 was analyzed using a new mathematical method. This method, called Logical Analysis of Data (LAD), is based on combinatorics, optimization and the theory of Boolean functions, and was adapted to handle the type of "inseparable" data which are typical for medical applications. LAD was introduced in 1986 [9], the first paper on the topic appeared in 1988 [6], and since then numerous papers (e.g., [1,2,7,8]) have been devoted to various mathematical, computational and applied aspects of this method.

* A preliminary version of this paper was presented at the *Annual Scientific Session of the American College of Cardiology* – March 18–21, 2001, Orlando, FL; see also [12].

The objective of this study was the construction of a model for distinguishing groups of patients at high and at low mortality risk. Risk stratification is a process common in medical practice by which patients are systematically assessed for the likelihood of developing a poor outcome [3]. For example, in cardiovascular medicine, patients with known or suspected coronary artery disease may undergo exercise testing with or without concurrent imaging in order to determine the risk of subsequent death or myocardial infarction [4,13–16,19]. The purpose of risk stratification is to identify high-risk patients who are most likely to benefit from aggressive therapy and low-risk patients who are best served by conservative care [3,19].

In the cardiovascular literature, risk stratification schemes are typically based on standard statistical models, such as logistic regression [10] or Cox proportional hazards [5]. A common problem with these approaches is that, although high-risk patients can be easily identified, they usually only account for a minority of subsequent clinical events [17]. Conversely, other risk markers may identify the majority of patients at high risk, but they also include in the same group sizeable numbers of other patients [11]. The ideal risk stratification scheme would identify a small subset of patients who will in fact account for the vast majority of deaths.

It will be shown that an appropriately modified version of LAD produced a model which classifies about 15% of the patients in a category having a mortality rate more than 4 times higher than the average, and about 70% of the patients in a category having a mortality rate of less than $1/5$ of the average. Moreover, $2/3$ of those patients who died during the observation period belong to the high-risk category defined by LAD.

A Prognostic Index, derived from the LAD model, is shown to have a 83.95% correlation with the mortality rate of patients. Using the Prognostic Index, the number of patients classified into the high-risk (20%) or low-risk (77%) categories was increased from a total of 85% to more than 97% of the population. Finally, it is shown that the classification given by the Prognostic Index agrees in 3 out of 4 cases with that of the Cox Score, widely used by cardiologists, and to outperform it slightly, but consistently. An example of a highly reliable risk stratification system using both indicators is described in section 10.

2. The problem and the data

The dataset consisted of observations about 9454 patients, 312 of whom died during the observation period. For each of the patients, 21 variables were recorded, including general data (age, gender), health history (chest pain, hypertension, diabetes, coronary artery disease), medication (beta blockers, verapamil, lipid lowering drugs, aspirin) and specific measurements (resting abnormal ECG, resting heart rate, change in heart rate, chronotropic index, duke treadmill score). Following [4,13,16], table 1 presents the complete list of the recorded variables. The analysis took into account all the variables, with the exception of #17 (“TTODEAD”), which was only used for cross validation. Variable 12 (“DEAD”), indicating whether the patient died during the observation period, was taken as the dependent variable.

Table 1
Attributes recorded for each patient.

	Attribute	Description
1	RESTST	resting abnormal ECG (0 = no, 1 = yes)
2	PT_AGE	age in years
3	RBBB	right bundle branch block on ECG (0 = no, 1 = yes)
4	BETABLOK	use of beta blockers (0 = no, 1 = yes)
5	DILVER	use of diltiazem or verapamil (0 = no, 1 = yes)
6	LIPIDRX	lipid lowering drugs (0 = no, 1 = yes)
7	COPD	chronic lung disease (0 = no, 1 = yes)
8	PVD	peripheral vascular disease (0 = no, 1 = yes)
9	GENDER	gender (0 = male, 1 = female)
10	RESTHR	resting heart rate in beats per minute
11	ASPIRIN	use of aspirin (0 = no, 1 = yes)
12	DEAD	died during followup (0 = no, 1 = yes)
13	CHESTP	history of chest pain (0 = no, 1 = yes)
14	SMKNOW	current smoker (0 = no, 1 = yes)
15	HTN	hypertension (0 = no, 1 = yes)
16	DIABETES	diabetes (0 = no, 1 = yes)
17	TTODEAD	length of follow-up in years
18	DHRREC	change in heart rate during the first minute of recovery in beats per minute
19	CRI	chronotropic index, a measure of heart rate rise during exercise
20	PRIORCAD	history of known coronary artery disease (0 = no, 1 = yes)
21	DUKE	duke treadmill score with typical values between -20 and $+15$

Table 2
Sample of observations.

Observation	RESTST	PT_AGE	RBBB	BETABLOK	DILVER	LIPIDRX	COPD	PVD	GENDER	RESTHR	ASPIRIN	CHESTP	SMKNOW	HTN	DIABETES	TTODEAD	DHRREC	CRI	PRIORCAD	DUKE	DEAD
1	0	65	0	0	0	0	0	0	0	84	0	0	0	1	0	6.5	27	1.20	0	8.5	0
2	0	68	0	0	0	0	0	0	1	77	0	0	0	0	0	8.1	13	1.04	0	9.5	0
3	0	63	0	0	0	0	0	0	0	60	0	0	0	0	0	8.4	29	1.01	0	11.5	0
4	1	70	0	0	1	0	0	0	0	83	1	0	0	0	1	6.7	8	1.15	1	8.5	0
5	0	68	0	0	0	0	0	0	0	71	0	0	0	0	0	8.1	8	0.91	0	6.5	0
6	0	60	0	0	1	0	0	0	1	89	1	0	0	0	0	8	24	0.75	1	9.5	0
7	0	74	0	0	0	0	0	0	0	85	1	0	0	1	0	7.5	21	1.38	1	9.5	0
8	0	79	0	1	1	0	0	0	1	69	0	0	0	1	0	7.5	6	0.47	1	3	0
9	1	73	0	0	1	1	0	0	1	89	1	0	0	1	0	8.5	15	0.78	1	-2	0
10	0	60	0	0	0	0	0	0	0	61	0	1	0	1	0	2.9	19	0.68	0	8.5	0
11	0	66	0	0	0	0	0	0	1	68	0	0	0	1	0	5.3	31	0.86	0	4.5	1
12	0	85	0	1	0	0	0	0	1	86	1	0	0	0	0	2.4	13	0.45	1	3	1
13	0	75	0	0	1	0	0	0	0	88	0	0	0	0	0	6.2	13	0.96	1	6.5	1
14	1	88	0	0	0	0	0	0	0	65	1	0	0	0	1	4.3	7	1.09	1	5.5	1

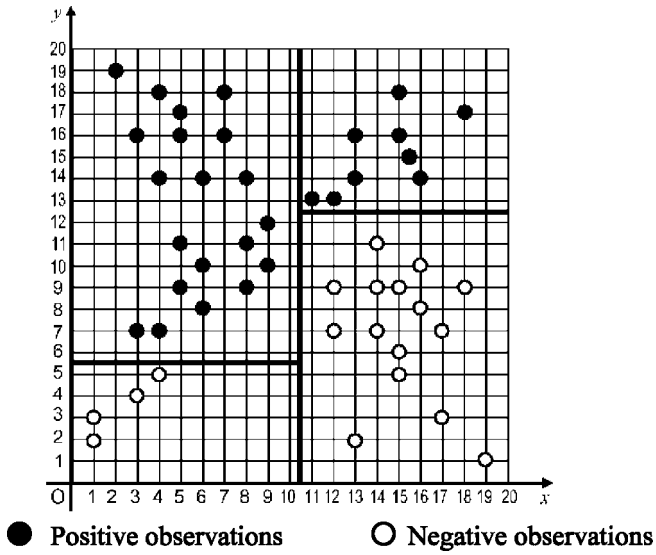

Figure 1. Example of separable dataset.

For illustration, we present in table 2 the data for ten of the surviving patients, and four of the patients who died during the observation period.

Table 2 points to a common situation in data analysis. It can be seen from the sample of these 14 observations that none of the variables can by itself make a distinction between those patients who died, and those who did not. Indeed, there is no binary variable taking one of the binary values for all the patients who died, and taking the opposite binary value for all those who survived. Similarly, there is no numerical variable taking only "small" values for one of the two groups of patients, and only "large" values for the other group. On the other hand, it will be seen that there are powerful, robust ways of distinguishing the two groups of data, and that the identification and use of such "classifiers" is the essential feature of LAD.

Let us introduce some terminology. In the discussion below we shall frequently identify a *patient* with an *observation*, and also with the *point in* $\mathbf{R}^{20}$, whose components are the corresponding values of the 20 *variables* (or *attributes*). Such a point will have a positive or a negative label, depending on the value of the *outcome* attribute "DEAD"; the observation will be labeled as *positive* if the patient died, and *negative* if he/she survived the observation period.

In previous applications of LAD and other data mining techniques, the $\mathbf{R}^n$ representation of the analyzed datasets generally admitted a more or less "crisp" separation into homogeneous zones, containing only positive or only negative points. An example of such a "separable" dataset is shown in figure 1, while figure 2 provides an example for an "inseparable" dataset. An obvious characteristic feature of the inseparable datasets consists in the fact that for many of the given data points are not contained in any "reasonably sized" homogeneous interval of $\mathbf{R}^n$ (i.e., interval containing only positive or only negative points); in the example of table 2 this property holds for the vast majority

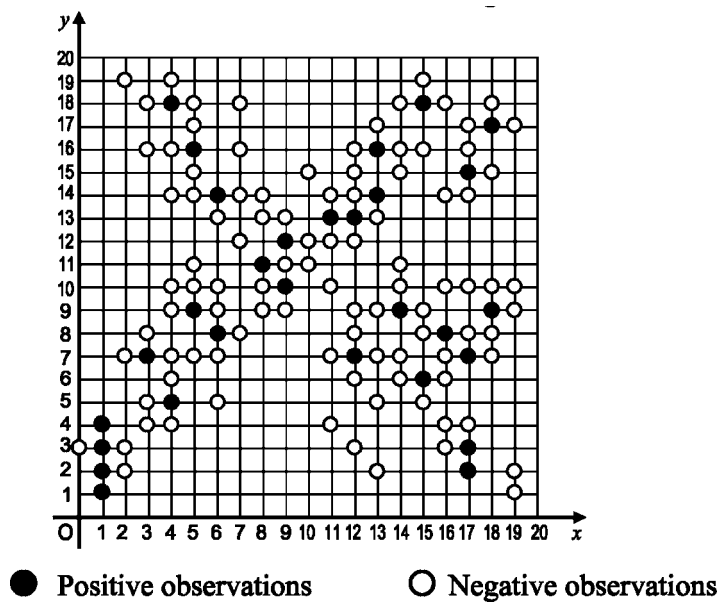

Figure 2. Example of inseparable dataset.

of the positive points. The inseparability of the positive and negative data is characteristic to many datasets occurring in medicine, finance and several other important areas of application.

In order to be able to handle inseparable data we have developed a modified version of the general LAD method, which will be described below and illustrated on the CCF dataset.

3. Patterns

Let us consider a dataset Ω consisting of a finite set of observations, represented as vectors in $\mathbf{R}^n$. We shall denote by Ω^+ and Ω^-, respectively, the subsets of positive and negative observations in Ω.

A subset of points in $\mathbf{R}^n$ identified by upper and/or lower bounds placed on some of their components will be called a *pattern*. More precisely, a *pattern* P is described with the help of two subsets of indices, I and J, and of a set of real numbers α_i $(i \in I)$ and β_j $(j \in J)$, called "cutpoints", and is defined as

$$P = \left\{x \in \mathbf{R}^n \mid x_i \leqslant \alpha_i \ (i \in I),\ x_j \geqslant \beta_j \ (j \in J),\ I, J \subseteq \{1, 2, \ldots, n\}\right\}. \qquad (1)$$

A pattern P will be called *positive* if $P \cap \Omega^+ \neq \varnothing$ and $P \cap \Omega^- = \varnothing$. Similarly, a pattern P will be called *negative* if $P \cap \Omega^- \neq \varnothing$ and $P \cap \Omega^+ = \varnothing$. Clearly, a positive pattern is an interval of $\mathbf{R}^n$ which contains some positive observations and no negative

ones. For example, the set of male patients aged 77 or older, having their resting heart rate of at least 97 beats/minute, i.e., those satisfying the conditions

$$\begin{cases} \text{PT\_AGE} \geqslant 77, \\ \text{GENDER} = 0, \\ \text{RETSTHR} \geqslant 97, \end{cases} \tag{2}$$

includes 5 positive observations from the CCF dataset and no negative observations, thus representing a positive pattern.

Positive and negative patterns, introduced in [6], represent in fact one of the fundamental tools used in LAD, and were seen in [2] to give a strong insight into the structure of numerous practical problems of data analysis.

Positive and negative patterns, in their original "pure" form, represent homogeneous intervals in $\mathbf{R}^n$, containing only positive or only negative points. In analyzing the CCF dataset it can be seen that the concepts of positive and negative patterns will have to be extended. One reason for this is the fact that this dataset is quite unbalanced, containing 3.3% positive points (representing patients who died during the observation period), and 96.7% negative points (representing surviving patients). Moreover, the relatively small set of points in $\mathbf{R}^{20}$ representing positive observations is dispersed among the much larger set of points representing negative observations. The presence of *inseparable data* is strongly related to the non-deterministic nature of the relation between the dependent variable and the values of the independent variables. While it is obvious that the fulfillment of various conditions can substantially increase or decrease the *risk of death*, it is equally obvious that the fulfillment of these conditions cannot *guarantee* the occurrence or the absence of this event during the observation period.

As a consequence of the inseparable character of the data, the only positive and negative patterns are of a small size, i.e., contain few data points. For example, the "richest" positive pattern detected in the CCF dataset contains only 11 positive data points, and most of the remaining positive patterns contain not more than 4 or 5 points. Clearly, no conclusions based on such "small" positive patterns can carry much weight. In order to handle this problem, we shall generalize the concept of patterns, allowing beside "pure" positive and negative patterns, also those for which the proportion between the positive and negative observations contained in them satisfy some conditions. The requirements to be imposed on these "generalized" patterns (which, for the sake of simplicity, will be simply called "patterns") will be described in the next section.

4. Characteristics of patterns

The number of patterns which can be detected in a dataset is extremely large. For example, the number of those patterns which can be described using at most three variables in the CCF dataset exceeds $29 \cdot 10^6$. It is therefore important to restrict the attention to reasonably sized subsets of patterns having high information value. In order to extract such subsets, we shall associate to every pattern several parameters, and will limit our search

only to those patterns whose associated parameters are below or above some threshold values.

The basic criteria for the experimental selection of parameter-bounds are the following. First, the bounds have to be such that the resulting patterns should be able to provide a powerful distinction between the classes of positive and negative observations. Second, the collection of patterns satisfying the requirements should "cover" the dataset, i.e., each (or at least most of the) observation point(s) should satisfy at least some of the patterns in the collection. Third, the patterns in the collection should be "rich", i.e., the average number of observations satisfying a pattern in the collection should not be too small. Fourth, whenever possible, the resulting patterns should be easily "understandable", i.e., the average number of variables appearing in their expressions ($|I| + |J|$ in the notations of (1)) should be small.

While the parameters can be defined in a rigorous way, and the criteria for imposing limitations on them can also be clearly formulated, the "fine-tuning" of the parameter-bounds is carried out through experimentation.

The three basic parameters to be used in this paper are the concepts of *risk*, *degree* and *prevalence* associated to a somewhat generalized concept of patterns. These concepts along with the associated thresholds used in this analysis will be presented below.

4.1. Risk

In order to analyze problems with inseparable data we shall first of all define the concept of *risk* ρ_P of a pattern P as the proportion of positive points in the pattern

$$\rho_P = \frac{|P \cap \Omega^+|}{|P \cap \Omega|}; \tag{3}$$

here $|X|$ represents the cardinality of the set X. We shall also introduce some real numbers t^+ and t^-, called the *high-risk* and the *low-risk thresholds*, respectively, and shall say that a pattern is of *high risk* (or of *low risk*) if its risk is at least t^+ (respectively, at most t^-).

Consider for example, the patterns P_1 and P_2 defined by

$$P_1 = \begin{cases} \text{RESTST} = 0, \\ \text{PT\_AGE} \geqslant 58, \\ \text{RESTHR} \geqslant 94 \end{cases} \tag{4}$$

and

$$P_2 = \begin{cases} \text{RESTST} = 1, \\ \text{PT\_AGE} \geqslant 64, \\ \text{GENDER} = 0. \end{cases} \tag{5}$$

The defining conditions of pattern P_1 are satisfied by 168 observations, 33 of which are positive. Similarly, the defining conditions of pattern P_2 are satisfied by 347 observations, 54 of which are positive. Clearly, $\rho_{P_1} = 19.64\%$, and $\rho_{P_2} = 15.56\%$. If, for

example, we decide to choose 16.5% as the value of t^+, then P_1 will be viewed as a high-risk pattern, while P_2 will not.

In the case of the CCF dataset, we have identified all patterns for which the mortality rate of the patients satisfying it is at least 5 times larger than the mortality rate of the entire observed population (i.e., it is at least $5 \cdot 3.3\%$); therefore we took $t^+ = 16.5\%$. Similarly, we have identified all patterns for which the mortality rate of the patients satisfying it is at most $1/5$ of the mortality rate of the entire observed population (i.e., it is at most $1/5 \cdot 3.3\%$); therefore we took $t^- = 0.66\%$.

For further illustration, the pattern

$$P_3 = \begin{cases} \text{PT\_AGE} \geqslant 68, \\ \text{CHESTP} = 0, \\ \text{DUKE} \leqslant 5.5 \end{cases} \tag{6}$$

contains 348 patients, of which 59 died during the observation period, i.e., $\rho_{P_3} = 16.95\%$, and is therefore a high-risk pattern. Similarly, the pattern

$$P_4 = \begin{cases} 50 \leqslant \text{PT\_AGE} \leqslant 52, \\ \text{CRI} \geqslant 1.05 \end{cases} \tag{7}$$

contains 955 patients, of which 5 died during the observation period, i.e., $\rho_{P_4} = 0.55\%$, and is therefore a low-risk pattern.

4.2. Degree

The degree of a pattern is the number of inequalities appearing in its definition (1). For example, both the high-risk pattern P_3 and the low-risk pattern P_4 contain 3 inequalities in their definition. Consequently, each of them is of degree 3.

Usually, we shall concentrate on patterns of low degree, due to two reasons: first their meaning can be easily understood by field experts, and second they can be produced without major computational difficulties.

The analysis of the CCF dataset was based on the use of the high-risk and low-risk patterns of degree at most 3. It should be remarked however, that some high-degree patterns might have a particularly strong information content. For example, the pattern

$$P_5 = \begin{cases} \text{PT\_AGE} \leqslant 59, \\ \text{RESTHR} \leqslant 126, \\ \text{DHRREC} \geqslant 12, \\ \text{DUKE} \geqslant -3.5 \end{cases} \tag{8}$$

of degree 4 contains the remarkable number of 5915 negative points (64.70% of all negative points in the dataset), and contains only 46 positive points, implying that the death rate of patients whose data satisfy this pattern is of only 0.77%, i.e., less then a quarter of the 3.3% mortality rate of the entire population.

Our experiments with the CCF dataset have shown that in spite of the existence of some higher-degree patterns of remarkably powerful characteristics, confining the

analysis to patterns of degree at most 3 leads to results comparable to the analysis of all patterns of degree 4, or even 5. The analysis of all patterns of degree 6 or higher in our dataset is computationally expensive, and not likely to lead to a substantial improvement of the results.

4.3. Prevalence

The number of observations contained in a pattern is called its *absolute prevalence*. The *absolute positive* or *negative prevalence* of a pattern is the number of positive or negative observations contained in it.

In a similar way, the ratio between the absolute prevalence of a pattern and $|\Omega|$ is called the *relative prevalence* of the pattern, while the ratio between the absolute positive (negative) prevalence and $|\Omega^+|$ (respectively $|\Omega^-|$) is called the *relative positive* (respectively *negative*) *prevalence* of the pattern; clearly the absolute negative (positive) prevalence of a positive (respectively negative) pattern is zero.

For illustration, the relative prevalence of the pattern P_4, defined by (7) is $955/9454$, i.e., 10.1%; its relative negative prevalence is $950/9142$, i.e., 10.4%, while its relative positive prevalence is $5/312$, i.e., 1.6%.

The information content of a pattern with a low relative positive and negative prevalence is minimal. Therefore, a significant analysis must be restricted only to patterns having a sufficiently high relative positive or negative prevalence. In this study we shall restrict our attention only to patterns whose relative positive or negative prevalences are of at least 10%.

To summarize, in this study we shall restrict our attention to those high-risk and low-risk patterns, which satisfy the following 3 conditions:

(i) have degree at most 3,

(ii) are defined by high- and low-risk thresholds of 16.5% and 0.66%, respectively,

(iii) have relative positive or negative prevalences of at least 10%.

5. Cutpoints

It has been seen before that patterns are defined by a system of simple inequalities of the form (1). The values of α_i and β_j appearing on the right hand sides of the inequalities (1) are called *cutpoints*.

Originally [6], LAD was developed for the analysis of binary data. When the method was extended to the analysis of numerical data, cutpoints were introduced [1] in order to replace each numerical variable x by several binary variables x_k, which were defined as taking the value 1 if the value of x exceeded some value s_k, and taking the value 0 otherwise. The cutpoints s_k were determined in such a way as to separate as well as possible the positive observations from the negative ones. Although, binary variables do not need cutpoints, for the sake of uniformity we can consider 0.5 to be their only

cutpoint. The minimization of the total number of cutpoints to be used was formulated in [1] as a set-covering problem, and it was shown there that most of the natural questions concerning this parameter turn out to be NP-hard, even in the case of only 2 numerical variables.

The determination of cutpoints in the logical analysis of numerical data was based on the assumption that the patterns defined by them will either be positive or negative. In the case of the CCF dataset, as well as in many of the medical, financial, etc datasets, the inseparability of the data renders any attempt to cover the data space with positive and negative patterns of sufficiently large prevalence, futile. Therefore, the logical analysis of this type of datasets has to be based on the use of high-risk and low-risk patterns, rather than positive and negative ones.

A most natural way to identify the cutpoints to be used in the definition of high-risk and low-risk patterns is to use an "equipartitioning" system, which can be built in the following way. Let us assume that for a given dataset the values taken by one of the variables, say x, belong to a set $V = \{v_1, v_2, \ldots, v_m\}$, where each of the values v_i actually occurs in the dataset, and where $v_1 < v_2 < \cdots < v_m$. Let further s_i be the midpoint of the interval $[v_i, v_{i+1}]$. The only inequalities $x \leqslant s$ or $x \geqslant s$ which will be used in the definition of patterns are those corresponding to the values $s_1, s_2, \ldots, s_{m-1}$ defined above. The validity of restricting our attention only to these values has been noticed in [1], where additional simplifications were also introduced for the elimination from consideration of some of these values s_i.

In order to restrict further the number of cutpoints in the system $s_1, s_2, \ldots, s_{m-1}$, we extract from it a subsystem of values $s_{i_1}, s_{i_2}, \ldots, s_{i_p}$ in such a way that the number of those observations whose component x takes values in anyone of the intervals $[s_q, s_{q+1}]$ should be approximately the same for all $q = 1, 2, \ldots, p$.

In the case of the CCF dataset, after comparing the results for $p = 10, 20, 30, 40$, we found the quality of the results to be highest for $p = 30$, and used therefore the corresponding equipartitioning system for the definition of all patterns considered in this study. For illustration we list in table 3 the cutpoints for the five numerical variables of the CCF dataset.

6. The pandect

The proposed method requires the selection of a set of cutpoints and the generation of all patterns having prescribed bounds on their degrees, their relative prevalences, and their risk. As mentioned in section 4, in the case of the CCF dataset, these "parameters" were chosen in the following way: the degrees were required not to exceed 3, both the positive and negative relative prevalences were 10%, or higher, and the high- and the low-risk thresholds were taken respectively as 16.5% (i.e., five times the mortality rate of the entire population), and 0.66% (i.e., one fifth of the mortality rate of the entire population).

It should be remarked that the collections of intervals defined by other choices of the parameters could be quite different. The actual identification of informative values of

Table 3
Cutpoints for the 5 numerical variables.

Cutpoint	Numerical variables				
	PT_AGE	RESTHR	DHRREC	CRI	DUKE
1	44.0	58.0	5.0	0.541	−1.5
2	47.0	61.5	8.0	0.614	0.5
3	48.0	63.0	9.0	0.637	1.3
4	49.0	64.5	10.0	0.666	2.3
5	50.0	65.5	10.5	0.704	3.3
6	51.0	67.0	11.5	0.729	3.8
7	52.5	68.5	12.0	0.751	4.0
8	53.5	69.5	13.0	0.771	4.3
9	54.0	70.5	13.5	0.797	4.8
10	55.0	71.0	14.0	0.812	5.0
11	55.5	72.0	14.5	0.826	5.5
12	56.0	73.0	15.0	0.839	5.8
13	57.0	74.0	15.5	0.850	6.0
14	58.0	74.5	16.0	0.863	6.3
15	58.5	75.5	16.5	0.876	6.5
16	59.0	76.4	17.0	0.885	6.5
17	60.0	76.5	17.5	0.898	7.0
18	61.0	78.0	18.0	0.911	7.0
19	61.5	79.0	18.5	0.932	7.5
20	62.0	79.5	19.0	0.940	7.5
21	62.5	80.5	19.5	0.956	7.8
22	63.5	81.5	20.0	0.968	8.0
23	64.5	82.5	21.0	0.984	8.3
24	65.0	84.0	21.5	1.002	8.5
25	66.5	85.5	22.5	1.020	8.8
26	67.5	86.5	23.0	1.039	9.0
27	68.5	88.0	23.5	1.061	9.5
28	70.0	90.5	25.0	1.081	10.0
29	71.5	92.5	26.0	1.104	10.5
30	74.0	98.0	28.0	1.152	11.0

the parameters is an experimental process consisting in the creation of various collections of such patterns defined by various parameter values, and the selection of those with the most satisfactory properties.

Taking into account the above definition of the parameters, and using the cutpoint system described in section 5, we have generated all the (approximately $29 \cdot 10^6$) patterns of degree at most 3, and selected from them all those patterns which satisfy our conditions concerning their prevalences and risks. The computer time needed to perform this operation using an Intel III/1 GHz processor was of approximately 450 seconds.

Let us denote by Σ the set of all high-risk and low-risk patterns satisfying the prescribed conditions. Σ will be called the *pandect* defined by the chosen system of cutpoints and parameter values. Let further $\Sigma = \Sigma^+ \cup \Sigma^-$, where Σ^+ and Σ^- are the subsets of high-risk and low-risk patterns, respectively.

The pandect Σ constructed for the CCF dataset using the parameter values described before (30 equipartitioning cutpoints/numerical variable, pattern degrees $\leqslant 3$, $t^+ = 16.5\%$, $t^- = 0.66\%$, relative positive and negative prevalences of at least 10%) consisted of 783 high-risk and 3940 low-risk patterns.

In order to introduce a quality measure of the pandect, or of any subset S of patterns, with $S = S^+ \cup S^-$, where $S^+ \subseteq \Sigma^+$ and $S^- \subseteq \Sigma^-$, let us define the *explained dataset* E_S as being the set consisting of those elements of Ω which satisfy at least one of the patterns in S. Similarly, let us define the *unexplained dataset* U_S as being $\Omega - E_S$. One of the quality measures of a system S of patterns is the proportion $u_S = |U_S|/|\Omega|$; if the system "explains" the whole set Ω then u_S is zero.

A second important quality measure of S is the *overlap* ε_S, defined as

$$\varepsilon_S = \frac{1}{|\Omega|}\left|\left(\bigcup_{\sigma\in S^+}\sigma\right)\cap\left(\bigcup_{\tau\in S^-}\tau\right)\right|. \tag{9}$$

Clearly, the overlap of S indicates the proportion of those points of Ω which satisfy some of the high-risk, as well as some of the low-risk patterns; for an ideal system ε_S should be zero.

7. Theories and models

Clearly, the pandect Σ, containing 4723 high- and low-risk patterns is much too large for practical applications. Moreover, some of the positive and negative observations may be contained in high- or low-risk patterns in Σ, indicating a large redundancy of the pandect. It is natural therefore to investigate nonredundant subsystems of patterns. Obviously, the explained dataset $E_{\Sigma'}$ of such a subsystem cannot be larger than the explained dataset E_Σ of the complete set Σ.

We shall call a subset $\mathbf{T}^+$ of high-risk patterns of Σ^+ a *positive theory* if $E_{\mathbf{T}^+} = E_{\Sigma^+}$, i.e.,

$$x \in \Omega^+ \cap \left(\bigcup_{\sigma\in\Sigma^+}\sigma\right) \Rightarrow x \in \bigcup_{\sigma\in\mathbf{T}^+}\sigma; \tag{10}$$

clearly, every observation contained in some high-risk pattern of the pandect is also contained in a high-risk pattern of any positive theory.

Similarly, we shall call a subset $\mathbf{T}^-$ of low-risk patterns of Σ^- a *negative theory* if $E_{\mathbf{T}^-} = E_{\Sigma^-}$, i.e.,

$$x \in \Omega^- \cap \left(\bigcup_{\sigma\in\Sigma^-}\sigma\right) \Rightarrow x \in \bigcup_{\sigma\in\mathbf{T}^-}\sigma; \tag{11}$$

clearly, every observation contained in some low-risk pattern of the pandect is also contained in a low-risk pattern of any negative theory.

A pair consisting of a positive and a negative theory will be called a *model*.

It is clear that the minimum sizes of the positive and of the negative theories can be determined by solving separately two naturally associated set-covering problems.

In the case of the CCF dataset we have determined a model **T** consisting of 42 high-risk and 77 low-risk patterns. The model is presented in tables 4 and 5, where the prevalences and the risks of each pattern are specified, along with the inequalities defining the pattern in terms of the original variables.

The quality of the model can be measured by the size of its unexplained dataset $U_{\mathbf{T}}$, which consists of 247 points (i.e., $u_{\mathbf{T}} = 2.61\%$), and of its overlap $\varepsilon_{\mathbf{T}} = 12.26\%$.

In order to get a more detailed understanding of the classification power of the model, we shall introduce now the concept of a classification matrix. The *classification matrix* of a model **T** consists of 4 rows (corresponding respectively to positive data (P), negative data (N), *Risk*, and prevalence (*Prev*)), and 4 columns:

1. *HRT* – percentage of observations satisfying high-risk patterns of **T** only,
2. *LRT* – percentage of observations satisfying low-risk patterns of **T** only,
3. *HLT* – percentage of observations satisfying both high and low-risk patterns of **T**,
4. *UT* – percentage of observations not satisfying any of the high or low-risk patterns of **T**.

In an ideal case, the only nonzero element in the row P would be the entry 100% appearing in column *HRT*, and similarly the only nonzero entry in the row N would be 100% appearing in the column *LRT*. In the same case the first two entries in the row *Risk* would be 100% and 0%, respectively, and the only two nonzero entries in the row *Prev* would appear in the first two entries and would add up to 100%.

For the CCF dataset the model **T** produces the classification matrix presented in table 4. As an example we show how the risk for the column *HRT* in table 4 was calculated. This column includes 63.78% of the 312 positive observations representing patients who died during the observation period (i.e., 199 patients), and 13.17% of the 9142 negative observations (i.e., 1204 patients). Therefore, the risk in the high-risk subset of observations reported in this column is $199/(199 + 1204)$, i.e., 14.18%.

In conclusion, the model:

- identifies a set of high-risk observations with a mortality rate of 14.18%, and a set of low-risk observations with a mortality rate of 0.65%,

Table 4
Classification matrix for model.

	HRT	*LRT*	*HLT*	*UT*
P	63.78%	13.78%	20.83%	1.60%
N	13.17%	72.22%	11.97%	2.65%
Risk	14.18%	0.65%		
Prev	14.84%	70.28%	12.26%	2.61%

Table 5
Positive theory consisting of 42 high-risk patterns.

	Prevalence		Pattern parameters			Pattern description											
Pattern	Positive	Negative	Degree	Relative positive prevalence	Risk	RESTST	PT_AGE	PVD	GENDER	RESTHR	ASPIRIN	RESTCP	HTN	DIABETES	DHRREC	CRI	PRIORCAD	DUKE
P1	38	178	3	12.18%	17.59%	1	>65											
P2	50	242	3	16.03%	17.12%	1	>60.5									<0.823		
P3	39	191	3	12.50%	16.96%	1	>62											<4
P4	35	171	3	11.22%	16.99%	0										<0.719	1	
P5	40	198	3	12.82%	16.81%		>65	0										<1.5
P6	49	241	3	15.71%	16.90%		>58		0	>88								
P7	85	422	3	27.24%	16.77%		>56.5		0							<0.762		
P8	46	224	2	14.74%	17.04%		>65			>86								
P9	39	181	3	12.50%	17.73%		>58			>88			1					
P10	45	227	3	14.42%	16.54%		>42.5			>93.5					<11			
P11	72	361	3	23.08%	16.63%		>55.5			>72					<9			
P12	61	301	3	19.55%	16.85%		>55.5			>83					<11			
P13	79	394	3	25.32%	16.70%		>58			>61.5					<9			
P14	58	280	3	18.59%	17.16%		>60.5			>83					<13.5			
P15	58	287	3	18.59%	16.81%		>65			>76.5					<14.5			
P16	45	222	3	14.42%	16.85%		>60.5			>88						<1.119		
P17	50	250	3	16.03%	16.67%		>63.5			>76.5						<0.924		
P18	56	283	3	17.95%	16.52%		>65			>78.5						<1.097		
P19	46	232	3	14.74%	16.55%		>63.5			>74.5								<4
P20	58	289	3	18.59%	16.71%		>65			>76.5								<6.25
P21	68	341	3	21.79%	16.63%		>53				0					< 0.762		
P22	55	272	3	17.63%	16.82%		>65				0					< 0.877		
P23	72	359	3	23.08%	16.71%		>65						1			<1.008		
P24	69	349	3	22.12%	16.51%		>52.5								<9	<0.891		
P25	33	167	3	10.58%	16.50%		>58								<14.5			<1.5
P26	94	452	3	30.13%	17.22%		>62								<11			<9
P27	71	358	3	22.76%	16.55%		>65								<22			<4
P28	87	438	3	27.88%	16.57%		>67.5								<14.5			<11
P29	82	412	3	26.28%	16.60%		>67.5								<17.5			<5.5
P30	59	296	3	18.91%	16.62%		>67.5								<26.5			<3
P31	63	317	3	20.19%	16.58%		>55.5									<0.762	1	
P32	39	196	3	12.50%	16.60%		>59									<0.924		<1.5
P33	74	373	3	23.72%	16.55%		>62									<0.97		<3
P34	65	323	3	20.83%	16.75%				0						<12			<4.5
P35	45	225	3	14.42%	16.67%					>88					<11	<0.953		
P36	43	208	3	13.78%	17.13%						0				<16.5		1	
P37	48	241	3	15.38%	16.61%						0					<0.659		<11
P38	59	296	3	18.91%	16.62%						0					<0.796		<4.5
P39	32	161	3	10.26%	16.58%						0						1	<4.5
P40	87	416	3	27.88%	17.30%							0			<9			<7.5
P41	37	180	3	11.86%	17.05%										<16.5	<0.944		<1.5
P42	35	174	2	11.22%	16.75%											<0.762		<1.5

Table 6
Positive theory consisting of 77 low-risk patterns.

	Prevalence		Pattern parameters			Pattern description											
Pattern	Positive	Negative	Degree	Relative negative prevalence	Risk	RESTST	PT_AGE	PVD	GENDER	RESTHR	ASPIRIN	RESTCP	HTN	DIABETES	DHRREC	CRI	PRIORCAD	DUKE
N1	31	4709	2	51.51%	0.65%	0	<54											
N2	20	3132	1	34.26%	0.63%		<47.5											
N3	7	1140	2	12.47%	0.61%		<57.5			<63								
N4	5	935	2	10.23%	0.53%		<58			<61.5								
N5	6	948	3	10.37%	0.63%		>46, <52			<73								
N6	6	915	3	10.01%	0.65%		>48, <55			<71.5								
N7	10	1561	2	17.08%	0.64%		<60								>26.5			
N8	20	3062	2	33.49%	0.65%		<63.5									>1.008		
N9	5	921	3	10.07%	0.54%		>51.5, <56.5									>0.924		
N10	30	4516	2	49.40%	0.66%		<52										0	
N11	6	975	3	10.67%	0.61%		>49.5, <53											>7
N12	6	944	3	10.33%	0.63%		<56.5		1						>14.5			
N13	4	942	3	10.30%	0.42%		<53		1							>0.762		
N14	3	928	3	10.15%	0.32%		<57.5		1							>0.862		
N15	4	919	3	10.05%	0.43%		<60		1							>0.902		
N16	5	940	3	10.28%	0.53%		<63.5		1							>0.913		
N17	5	944	3	10.33%	0.53%		<55		1									>5
N18	6	920	3	10.06%	0.65%		<54			>74.5, <80								
N19	6	941	3	10.29%	0.63%		<55			>80, <88								
N20	6	926	3	10.13%	0.64%		<58			>64.5, <69								
N21	6	944	3	10.33%	0.63%		<59			<61.5				0				
N22	4	917	3	10.03%	0.43%		<60			<67					>24			
N23	6	921	3	10.07%	0.65%		<62			<71.5					>26.5			
N24	5	922	3	10.09%	0.54%		<57.5			>81					>17			
N25	4	931	3	10.18%	0.43%		<57.5			>82					>16			
N26	4	935	3	10.23%	0.43%		<58			>86					>12			
N27	5	929	3	10.16%	0.54%		<60			>79.5					>19.5			
N28	5	938	3	10.26%	0.53%		<60			>84.5					>15			

Table 6
(Continued).

	Prevalence		Pattern parameters			Pattern description											
Pattern	Positive	Negative	Degree	Relative negative prevalence	Risk	RESTST	PT_AGE	PVD	GENDER	RESTHR	ASPIRIN	RESTCP	HTN	DIABETES	DHRREC	CRI	PRIORCAD	DUKE
N29	6	923	3	10.10%	0.65%		<62			>82					>17			
N30	6	929	3	10.16%	0.64%		<63.5			>81					>18			
N31	5	953	3	10.42%	0.52%		<58			<63						>0.844		
N32	5	930	3	10.17%	0.53%		<60			<64.5						>0.902		
N33	6	935	3	10.23%	0.64%		<67.5			<67						>0.989		
N34	6	930	3	10.17%	0.64%		<59			>86						>0.924		
N35	6	938	3	10.26%	0.64%		<62			>86						>0.944		
N36	6	916	3	10.02%	0.65%		<65			>82						>1.038		
N37	5	933	3	10.21%	0.53%		<58			>86								>7
N38	5	915	3	10.01%	0.54%		<59								>22.5, <26.5			
N39	6	965	3	10.56%	0.62%		<54								>15			<7.65
N40	5	956	3	10.46%	0.52%		<56.5								>19.5			<9.5
N41	6	924	3	10.11%	0.65%		<59									>0.934, <0.98		
N42	6	928	3	10.15%	0.64%		<67.5									>1.05, <1.097		
N43	6	925	3	10.12%	0.64%		<55									>0.902		<8.5
N44	6	934	3	10.22%	0.64%		<60									>0.953		<7.65
N45	11	1837	2	20.09%	0.60%			0										>0
N46	5	963	3	10.53%	0.52%					>82	0							>10.9
N47	7	1116	3	12.21%	0.62%					<68			0					>0
N48	6	923	3	10.10%	0.65%					<67					>9	>0.998		
N49	6	958	3	10.48%	0.62%					<68					>15	>0.989		
N50	6	927	3	10.14%	0.64%					>74.5					>19.5	>0.989		
N51	6	973	3	10.64%	0.61%					>75					>18	>0.998		
N52	6	927	3	10.14%	0.64%					>75					>20.5	>0.934		
N53	6	925	3	10.12%	0.64%					>76					>19.5	>0.961		
N54	6	937	3	10.25%	0.64%					<69					>22.5			>10.9
N55	6	919	3	10.05%	0.65%					<69					>24.5			>5
N56	6	936	3	10.24%	0.64%					<73					>26.5			>6

Table 6
(Continued).

Pattern	Prevalence		Pattern parameters			Pattern description											
	Positive	Negative	Degree	Relative negative prevalence	Risk	RESTST	PT_AGE	PVD	GENDER	RESTHR	ASPIRIN	RESTCP	HTN	DIABETES	DHRREC	CRI	PRIORCAD	DUKE
N57	9	1417	3	15.50%	0.63%					>66					>21.5			>8.5
N58	9	1360	3	14.88%	0.66%					>71.5					>19.5			>7.65
N59	6	964	3	10.54%	0.62%					<70.5						>0.961, <1.05		
N60	6	932	3	10.19%	0.64%					<71.5						>0.97, <1.05		
N61	6	960	3	10.50%	0.62%					<72						>1.008, <1.097		
N62	6	927	3	10.14%	0.64%					<84.5						>0.998, <1.038		
N63	6	917	3	10.03%	0.65%					<69						>0.998		>6.5
N64	6	951	3	10.40%	0.63%					<83						>1.062		>8.5
N65	11	1673	3	18.30%	0.65%					>64.5						<1.097		>0
N66	5	937	3	10.25%	0.53%					>76						>1.028		>10.9
N67	6	933	3	10.21%	0.64%					>79.5						>1.018		>8.5
N68	5	926	3	10.13%	0.54%					>83							0	>10.9
N69	7	1084	3	11.86%	0.64%										>21.5, <29.5	>0.98		
N70	12	1885	3	20.62%	0.63%										>19.5			>0
N71	6	946	3	10.35%	0.63%										>12	>0.989, <1.028		
N72	6	943	3	10.32%	0.63%										>14.5	>0.998, <1.038		
N73	6	956	3	10.46%	0.62%										>18	>0.989, <1.05		
N74	10	1527	3	16.70%	0.65%										>22.5	>0.862, <1.05		
N75	12	1859	3	20.33%	0.64%										>22.5	>0.924	0	
N76	6	932	3	10.19%	0.64%										>18	>1.062		>7
N77	6	940	3	10.28%	0.63%										>29.5	>0.823		>5

- classifies correctly and unambiguously 63.78% of the positive observations as being at high-risk and 72.22% of the negative ones as being at low-risk,
- classifies erroneously 13.78% of the positive observations and 13.17% of the negative ones,
- does not provide any classification for 2.61% of the observations.

8. Prognostic Index

The classification provided by a model assigns a risk category for those patients who satisfy only high-risk or low-risk patterns.

We shall introduce in the this section a classification tool, which refines the classification provided by a model, by assigning a *Prognostic Index* to each patient, thus making it possible to assign a risk category also to those patients who satisfy both high-risk and low-risk patterns.

Let τ^+ be the number of high-risk patterns in the model $\mathbf{T}$, and $\tau^+(x)$ be the number of those high-risk patterns which are satisfied by a particular observation x; the indices τ^- and $\tau^-(x)$ are defined in a similar way.

The *Prognostic Index* $\pi(x)$ of patient x is defined as

$$\pi(x) = \frac{\tau^+(x)}{\tau^+} - \frac{\tau^-(x)}{\tau^-}. \tag{12}$$

Clearly, $\pi(x)$ will have a positive value for observations satisfying only high-risk patterns, and a negative value for those satisfying only low-risk patterns. By classifying now an observation x as being at high-risk or at low-risk according to the positive or negative sign of $\pi(x)$, we shall refine the classification provided in section 7 (since the classification defined in that section will never be contradicted).

The classification matrix for the Prognostic Index is defined very similarly to that in section 7, noticing simply that only 3 columns (*HRI*, *LRI* and *UI*) are needed this time, since the set *HLI* of patients having both positive and negative indices is obviously empty. For the CCF dataset the Prognostic Index π produces the classification matrix presented in table 7.

Clearly, the high-risk set defined by the Prognostic Index contains 1920 patients, i.e., 20.31% of the entire population, and has a mortality rate of 12.14%. The low-risk

Table 7
Classification matrix for Prognostic Index.

	HRI	*LRI*	*UI*
P	74.68%	23.72%	1.60%
N	18.45%	78.89%	2.66%
Risk	12.14%	1.02%	
Prev	20.31%	77.07%	2.63%

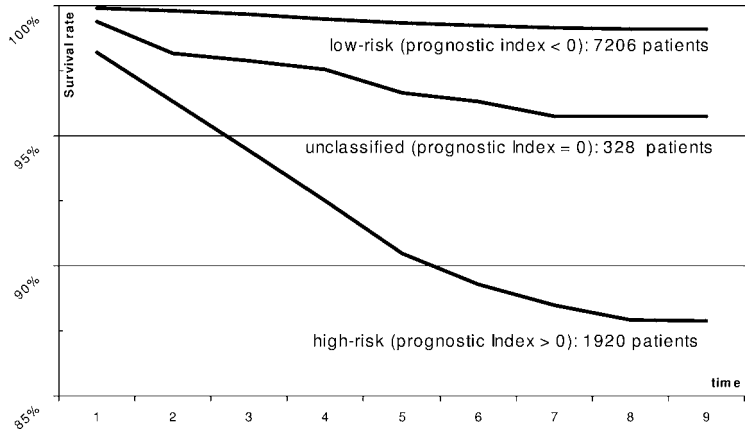

Figure 3. Dynamics of survival rates for low-risk, unclassified and high-risk patterns.

set contains 7286 patients, i.e., 77.07% of the entire population, and has a mortality rate of 1.02%.

In order to establish a clear relationship between the Prognostic Index values with mortality rates, we have re-indexed the patients in the decreasing order of their Prognostic Indices, and calculated, the average Prognostic Index API_i and the mortality rate MR_i of the groups of patients ranked 1–200, 201–400, 401–600, ... The *high correlation*, 83.95%, between these two sequences shows clearly that the Prognostic Index is an highly indicative estimator of the risk of death.

The graphs in figure 3 present the actual changes in survival rates of the three index-defined categories of patients (low-risk, unclassified and high-risk) along the nine years of observation, clearly confirming the above conclusions, and illustrating the sharp differences between categories.

It can be seen that the average year-by-year survival rate among patients classified as low-risk is of 99.9%, while that among patients classified as high-risk is 98.6%.

9. Model validation

In order to evaluate the accuracy of the model we have applied to it the standard 5-folding validation method, i.e., we have considered a random partition of the original dataset into 5 segments, each of them containing about one fifth of the positive and one fifth of the negative observations. Using this partition we have defined 5 LAD problems. In each of them one of the 5 segments of the original dataset was chosen to be used as "testing set" for validating the high-risk and low-risk patterns, as well as the positive and negative models obtained by applying LAD to the "training set", i.e., to the set consisting of the remaining 4 segments of the entire dataset Ω. For each of the 5 problems, using the chosen parameter values (30 cutpoints for each numerical variable, pattern degree at most 3, high-risk threshold 16.5%, low-risk threshold 0.66%, and positive and negative prevalences of at least 10% each) we have first calculated the risk and the relative

Table 8
Risk validation of pandect by 5-folding.

	Risk			
	High-risk		Low-risk	
Validation problem	Training	Testing	Training	Testing
1	16.82%	13.04%	0.41%	0.45%
2	16.88%	15.71%	0.41%	0.76%
3	16.86%	15.10%	0.40%	0.51%
4	16.82%	15.86%	0.39%	0.64%
5	16.84%	15.75%	0.40%	0.60%
Average	16.84%	15.09%	0.40%	0.59%

Table 9
Relative prevalence validation of pandect by 5-folding.

	Relative prevalence			
	Positive		Negative	
Validation problem	Training	Testing	Training	Testing
1	19.54%	17.56%	11.19%	10.76%
2	19.97%	17.06%	11.34%	11.51%
3	18.91%	17.45%	11.09%	11.16%
4	19.30%	17.95%	11.19%	10.81%
5	17.83%	16.07%	11.25%	11.62%
Average	19.11%	17.22%	11.21%	11.17%

prevalences of the pandects of high-risk and low-risk patterns of the 5 training sets, and recalculated them on the 5 corresponding testing sets.

The results of these calculations are reported in tables 8 and 9. It can be seen that the average risk among all the observations satisfying at least one of the high-risk patterns remains above 15% in the testing set, and that the average risk among all those observations which satisfy at least one of the low-risk patterns remains below 0.6%. Also, it can be seen that an average high-risk pattern is satisfied by more than 17% of the positive observations in the testing set, and the average low-risk pattern is satisfied by more than 11% of the negative observations.

Table 10 presents the results of applying the 5-folding validation method to the classification matrices corresponding to the prognostic index. It can be seen that the high-risk sets defined by this index have an average mortality rate of 10.75%, while the low-risk sets have an average mortality rate of only 1.3%.

10. Prognostic Index and Cox Score

The evaluations given by the Prognostic Index were compared to those given by the "Cox Scores" – a risk indicator widely used by cardiologists. It should be noted that

Table 10
Prognostic Index validation by 5-folding – classification matrices.

Validation problem		Training			Testing		
		HRI	*LRI*	*UI*	*HRI*	*LRI*	*UI*
1	*P*	78.00%	20.40%	1.60%	72.58%	24.19%	3.23%
	N	19.44%	78.78%	1.78%	20.35%	78.01%	1.64%
	Risk	12.06%	0.88%		10.79%	1.04%	
	Prev	21.37%	76.85%	1.77%	22.07%	76.23%	1.69%
2	*P*	77.20%	21.60%	1.20%	61.29%	38.71%	0.00%
	N	18.61%	79.63%	1.76%	17.18%	80.96%	1.86%
	Risk	12.42%	0.92%		10.80%	1.60%	
	Prev	20.54%	77.71%	1.75%	18.63%	79.57%	1.80%
3	*P*	78.40%	19.60%	2.00%	75.81%	22.58%	1.61%
	N	19.36%	78.70%	1.94%	20.51%	77.13%	2.35%
	Risk	12.16%	0.84%		11.14%	0.98%	
	Prev	21.31%	76.75%	1.94%	22.34%	75.33%	2.33%
4	*P*	79.60%	20.00%	0.40%	59.68%	35.48%	4.84%
	N	18.55%	79.82%	1.63%	18.93%	79.43%	1.64%
	Risk	12.79%	0.85%		9.66%	1.49%	
	Prev	20.57%	77.85%	1.59%	20.27%	77.98%	1.75%
5	*P*	75.00%	24.19%	0.81%	64.06%	32.81%	3.13%
	N	17.82%	80.88%	1.30%	17.49%	80.93%	1.58%
	Risk	12.49%	1.00%		11.36%	1.40%	
	Prev	19.71%	79.01%	1.28%	19.02%	79.34%	1.64%
		Average			Average		
		HRI	*LRI*	*UI*	*HRI*	*LRI*	*UI*
	P	77.64%	21.16%	1.20%	66.68%	30.75%	2.56%
	N	18.76%	79.56%	1.68%	18.89%	79.29%	1.81%
	Risk	12.38%	0.90%		10.75%	1.30%	
	Prev	20.70%	77.63%	1.67%	20.47%	77.69%	1.84%

in contrast with the Cox model which requires constant proportional hazards over time, LAD does not required the confirmation of any assumptions about the distribution of data or times to death events.

First of all, the Pearson correlation between the Prognostic Indices and the Cox Scores of the 9454 patients is 0.85, showing a high resemblance between the two indicators. In order to clarify whether a patient considered to be at high risk according to one of the indicators is also considered to be at high risk according to the other, we have considered the following measure of "agreement" between two indicators.

Let h be a number between 1 and 9454, and let U_h and V_h denote the sets of the h highest ranked patients according to the Prognostic Index and to the Cox Score, respectively. For every h we shall call the observations U_h and V_h to be at high h-risk, according to the Prognostic Index and Cox Score, respectively. Let us now define, for

Table 11
Agreement between Prognostic Index and Cox Score.

h	α_h	h	α_h	h	α_h	h	α_h
200	28.39%	2600	70.82%	5000	74.33%	7400	81.57%
400	40.99%	2800	72.19%	5200	74.40%	7600	82.58%
600	48.45%	3000	74.82%	5400	75.12%	7800	84.13%
800	50.61%	3200	74.57%	5600	75.16%	8000	85.08%
1000	55.12%	3400	75.07%	5800	75.09%	8200	86.57%
1200	56.43%	3600	75.13%	6000	75.77%	8400	88.13%
1400	60.07%	3800	75.27%	6200	76.58%	8600	89.66%
1600	61.68%	4000	74.10%	6400	77.45%	8800	91.47%
1800	64.29%	4200	74.09%	6600	78.11%	9000	93.55%
2000	65.96%	4400	73.91%	6800	78.64%	9200	95.81%
2200	67.16%	4600	74.20%	7000	79.39%	9400	99.07%
2400	70.38%	4800	73.97%	7200	80.72%	9454	100.00%

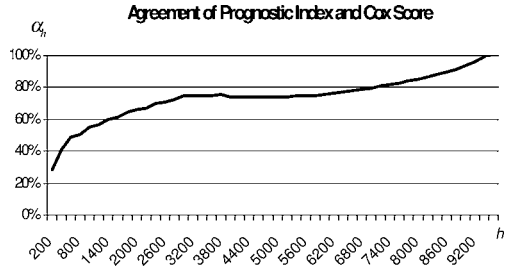

Figure 4. Agreement between Prognostic Index and Cox Score.

every h, the *agreement* α_h between U_h and V_h as

$$\alpha_h = \frac{|U_h \cap V_h|}{|U_h \cup V_h|}. \tag{13}$$

The values of α_h, calculated for $h = 200, 400, 600, \ldots$ are presented in table 11 and figure 4. The average value of α_h is 74%, giving further evidence to the high degree of agreement between the two indicators.

The comparison of the proportions of the 312 death events μ_h^{LAD} and μ_h^{COX} included in the classes U_h and V_h, is presented in table 12 and figure 5.

It can be seen from these results that the number of deaths among the top-ranked h patients is *consistently higher* if the ranking follows the Prognostic Index. Moreover, a 2-sided t-test shows that the probability for the proportion of deaths occurring in the group of h patients with highest Cox Scores to be equal to that in the group of h patients with highest Prognostic Indices is of $2.3 \cdot 10^{-6}$, and the probability to exceed it is $1.1 \cdot 10^{-6}$.

Further, repeating the same experiment for the groups of lowest ranked l patients according to the two indicators, we find the values given in table 13 and shown in figure 6.

Table 12
Proportion of deaths (out of 312) among h highest-ranked patients.

h	Prognostic Index	Cox Score
200	19.55%	16.35%
400	32.37%	29.81%
600	40.38%	38.14%
800	48.72%	44.55%
1000	56.09%	50.32%
1200	57.37%	57.05%
1400	61.86%	60.58%
1600	71.79%	63.14%
1800	74.04%	67.96%
2000	78.21%	71.15%
2200	78.21%	74.04%
2400	82.37%	75.96%
2600	82.37%	77.56%
2800	86.26%	80.13%
3000	88.14%	81.73%

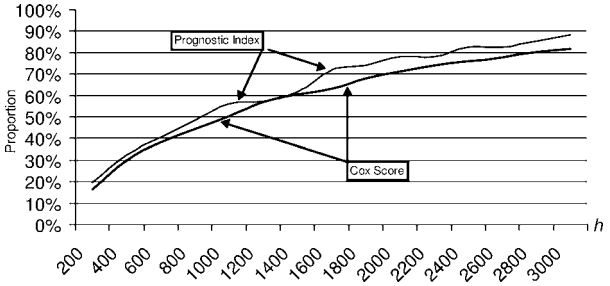

Figure 5. Proportion of deaths (out of 312) among h highest-ranked patients.

In agreement with the results in the previous experiment, these results show that the number of death events among the l lowest ranked patients is *consistently lower* when the ranking follows the Prognostic Index. Moreover, a 2-sided t-test shows that the probability for the proportion of deaths occurring in the group of l patients with lowest Cox Scores to be equal to that in the group of l patients with lowest Prognostic Indices is of $1.8 \cdot 10^{-3}$, and the probability to be smaller is of $3.6 \cdot 10^{-3}$.

Finally, let us examine the mortality rate within the groups of patients $U_h - V_h$ and $V_h - U_h$, i.e., the mortality rates in the groups of patients at high h-risk according to one of the indicators and at low h-risk according to the other. These rates are reported in table 14. In order to allow an immediate comparison with the mortality rate $\mu = 312/9454$ of the entire population, figure 7 expresses the results in table 14, as multiples of μ.

An obvious application of the above idea is to stratify the entire population taking into account both the Cox Score and the Prognostic Index. In order to simplify the stratification process, we divide the patients into quintiles in the decreasing order of

Table 13
Proportion of deaths (out of 312) among l lowest-ranked patients.

I	Prognostic Index	Cox Score
200	0.00%	0.26%
400	0.00%	0.38%
600	0.26%	0.51%
800	0.51%	0.64%
1000	0.77%	1.03%
1200	1.03%	1.28%
1400	1.03%	1.54%
1600	1.54%	1.92%
1800	2.44%	2.69%
2000	2.56%	3.08%
2200	3.33%	4.23%
2400	3.85%	5.64%
2600	4.74%	7.44%
2800	7.05%	9.49%
3000	8.72%	12.05%

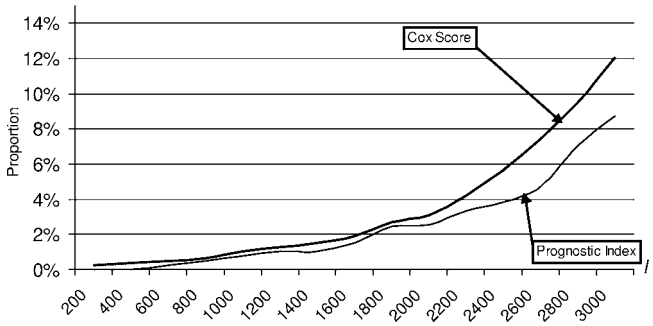

Figure 6. Proportion of deaths (out of 312) among l lowest-ranked patients.

their Cox Scores, and separately according to their Prognostic Indices. We shall say that a case has a level $PIL = k$, respectively $CSL = k$, if it is included in the kth quintile according to its Prognostic Index, respectively to its Cox Score.

In conclusion,

(i) the predictive value of the Prognostic Index resembles closely that of the Cox Score, with a small but consistent advantage to the former; moreover, whenever the high/low h-risk classifications provided by the two indicators differ, that one corresponding to the Prognostic Index is more informative;

(ii) in view of the striking difference in the underlying principles defining the two indicators, their close resemblance provides a strong validation for both of them;

(iii) the combined utilization of the two indicators provides a highly reliable risk stratification system.

Table 14

h	$U_h - V_h$	$V_h - U_h$	h	$U_h - V_h$	$V_h - U_h$
200	28.83%	19.82%	5000	1.36%	0.68%
400	17.96%	13.17%	5200	1.18%	0.66%
600	15.87%	12.50%	5400	0.91%	0.52%
800	13.36%	8.78%	5600	0.88%	0.50%
1000	12.11%	5.88%	5800	0.61%	0.61%
1200	7.19%	6.89%	6000	0.85%	0.36%
1400	7.45%	6.30%	6200	0.85%	0.36%
1600	10.29%	3.43%	6400	0.62%	0.37%
1800	8.70%	3.84%	6600	0.62%	0.49%
2000	7.80%	2.68%	6800	0.74%	0.49%
2200	5.56%	2.78%	7000	0.75%	0.50%
2400	6.24%	1.68%	7200	0.65%	0.39%
2600	5.18%	2.03%	7400	0.53%	0.40%
2800	4.65%	1.33%	7600	0.55%	0.28%
3000	5.56%	1.16%	7800	0.74%	0.30%
3200	4.29%	1.72%	8000	0.62%	0.31%
3400	4.55%	1.45%	8200	0.68%	0.17%
3600	3.72%	1.76%	8400	0.57%	0.19%
3800	3.36%	1.31%	8600	0.43%	0.00%
4000	2.69%	1.68%	8800	0.51%	0.00%
4200	1.92%	1.12%	9000	0.33%	0.00%
4400	1.82%	1.06%	9200	0.51%	0.00%
4600	1.62%	1.03%	9400	0.00%	0.00%
4800	1.25%	0.97%			

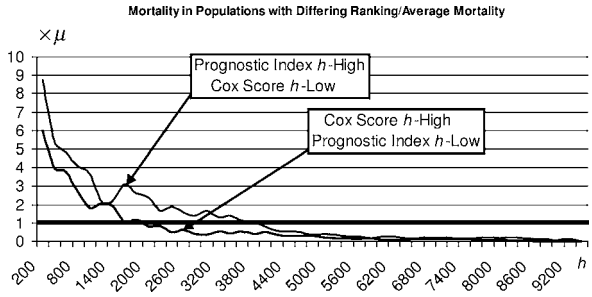

Figure 7.

11. Conclusions

Using Logical Analysis of Data on a population of patients under exercise stress testing, we have shown that it is possible to reliably identify a small subset of patients who are at relatively high risk of death, while simultaneously identifying a large population of patients who are at very low risk.

A characteristic feature of this dataset consists in the large disproportion between the number of the "positive" and of "negative" observations. This disproportion is due to

Table 15

Risk class	Cases contained	Class size	% of 312 deaths included	Risk
Very high	$PIL = CSL = 1$	1,483	63.78%	13.42%
High	$PIL,\ CSL \in \{1, 2\}$	3,222	85.58%	8.29%
Low	$PIL \in \{3, 4, 5\}$ *or* $CSL \in \{3, 4, 5\}$	6,232	14.10%	0.71%
Very low	$PIL,\ CSL \in \{4, 5\}$	2,957	2.56%	0.27%

the nature of the two classes of observations, which represent respectively the groups of patients who died or who survived during the observation period. While this disproportion is entirely reasonable in many medical (and some other) datasets, the methodological implications of it have not been previously examined in any other real-life application of LAD. From a purely methodological point of view this study led to an extension of the scope of LAD to the case of non-separable datasets.

A remarkable feature of the classification provided by LAD consists in the fact that the 20% segment of the population identified as being at high-risk included 3/4 of all those who died during the observation period; the mortality rate of this segment was 4 times higher than the mortality rate in the entire population. In contrast the 77% segment of segment of the population identified as being at low-risk had a mortality rate of less than 1/3 of the mortality rate in the entire population. Finally, the size of the unclassified segment of the population was of only 2.63%.

On the basis of the LAD model developed in this study, a Prognostic Index was defined for all patients, and its value was shown to be closely correlated with the patients' risk of death. A risk-of-death indicator widely used by cardiologist is the Cox Score. Comparing the Prognostic Index with the Cox Score, we have shown that their risk predictions coincide on the average in 3 out of 4 cases, and that the predictive value of the former outperforms slightly, but consistently, that one of the latter. The combined use of both indicators can make possible the construction of highly reliable risk stratification systems.

Most medical literature on risk stratification has focused on specific predictors of risk, with relatively less emphasis on interactions of risk factors, that is, on ways in which predictor variables affect each other's impacts on risk of an adverse outcome. While careful multivariable modeling makes it possible to examine two-way interactions, LAD, by its very nature, makes it possible to examine automatically tens of thousands of possible interactions with high degrees of complexity, retaining only the most significant ones. It can be expected that the interactions revealed through LAD may stimulate research for a better understanding of the related cause-effect relationships.

Among the useful features of LAD, we mention the possibilities it offers for explaining or justifying – on the basis of the patterns triggered by the values of the variables corresponding to a particular patient – the decision to classify the patient into a high-risk or a low-risk class, as well as decisions concerning the choice of particular alternatives for his or her treatment.

Acknowledgments

Drs. Lauer, Blackstone, and Ishwaran and Ms. Snader receive support from the National Heart, Lung, and Blood Institute (Grant RO1 HL-66004-1). Dr. Lauer, Dr. Blackstone, and Ms. Snader receive additional support from the American Heart Association (Established Investigator Grant 0040244N). Dr. Hammer and Mr. Alexe receive support from the National Science Foundation (Grant NSF-DMS-9806389) and the Office of Naval Research (Grant N00014-92-J-1375). We gratefully acknowledge all the support which made this study possible.

References

[1] E. Boros, P.L. Hammer, T. Ibaraki and A. Kogan, Logical analysis of numerical data, Mathematical Programming 79 (1997) 163–190.

[2] E. Boros, P.L. Hammer, T. Ibaraki, A. Kogan, E. Mayoraz and I. Muchnik, An implementation of logical analysis of data, IEEE Transactions on Knowledge and Data Engineering 12(2) (2000) 292–306.

[3] R.M. Califf, P.W. Armstrong, J.R. Carver, R.B. D'Agostino and W.E. Strauss, Stratification of patients into high, medium and low risk subgroups for purposes of risk factor management, J. Amer. Coll. Cardiol. 27(5) (1996) 1007–1019.

[4] C.R. Cole, E.H. Blackstone, F.J. Pashkow, C.E. Snader and M.S. Lauer, Heart-rate recovery immediately after exercise as a predictor of mortality, N. Engl. J. Med. 341(18) (1999) 1351–1357.

[5] D. Cox, Regression models and life tables (with discussion), J. Roy. Statist. Soc. Ser. B 34 (1972) 187–220.

[6] Y. Crama, P.L. Hammer and T. Ibaraki, Cause-effect relationships and partially defined Boolean functions, Annals of Operations Research 16 (1988) 299–326.

[7] O. Ekin, P.L. Hammer and A. Kogan, Convexity and logical analysis of data, Theoretical Computer Science 244 (2000) 95–116.

[8] P.L. Hammer, Y. Liu, B. Simeone and S. Szedmak, Saturated systems of homogeneous boxes and the logical analysis of numerical data, Annals of Operations Research (forthcoming), RUTCOR Research Report RRR 19-2000, RUTCOR, Rutgers University, Piscataway, NJ (2000).

[9] P.L. Hammer, Partially defined Boolean functions and cause-effect relationships, in: *International Conference on Multi-Attribute Decision Making via OR-Based Expert Systems*, University of Passau, Passau, Germany (1986).

[10] D. Hosmer and S. Lemeshow, *Applied Logistic Regression* (Wiley, New York, 1989).

[11] J.M. Kwok, T.D. Miller, T.F. Christian, D.O. Hodge and R.J. Gibbons, Prognostic value of a treadmill exercise score in symptomatic patients with nonspecific ST-T abnormalities on resting ECG, J. Amer. Medical Assoc. 282(11) (1999) 1047–1053.

[12] M.S. Lauer, S. Alexe, C.E. Pothier Snader, E.A. Blackstone, H. Ishwaran and P.L. Hammer, Use of the "Logical Analysis of Data" method for assessing long-term mortality risk after exercise electrocardiography, Circulation 106 (2002) 685–690.

[13] M.S. Lauer, G.S. Francis, P.M. Okin, F.J. Pashkow, C.E. Snader and T.H. Marwick, Impaired chronotropic response to exercise stress testing as a predictor of mortality, J. Amer. Medical Assoc. 282(6) (1999) 524–529.

[14] D.B. Mark, M.A. Hlatky, F.E. Harrell, Jr., K.L. Lee, R.M. Califf and D.B. Pryor, Exercise treadmill score for predicting prognosis in coronary artery disease, Ann. Intern. Med. 106(6) (1987) 793–800.

[15] D.B. Mark, L. Shaw, F.E. Harrell, Jr., M.A. Hlatky, K.L. Lee, J.R. Bengtson et al., Prognostic value of a treadmill exercise score in outpatients with suspected coronary artery disease [see comments], N. Engl. J. Med. 325(12) (1991) 849–853.

[16] T.H. Marwick, R. Mehta, K. Arheart and M.S. Lauer, Use of exercise echocardiography for prognostic evaluation of patients with known or suspected coronary artery disease, J. Amer. Coll. Cardiol. 30(1) (1997) 83–90.
[17] E.O. Nishime, C.R. Cole, E.H. Blackstone, F.J. Pashkow and M.S. Lauer, Heart rate recovery and treadmill exercise score as predictors of mortality in patients referred for exercise ECG, J. Amer. Medical Assoc. 284(11) (2000) 1392–1398.
[18] G. Rose, *The Strategy of Preventive Medicine* (Oxford University Press, New York, 1992).
[19] D.A. Weiner, T.J. Ryan, L. Parsons, L.D. Fisher, B.R. Chaitman, L.T. Sheffield et al., Long-term prognostic value of exercise testing in men and women from the coronary artery surgery study (CASS) registry, Amer. J. Cardiol. 75(14) (1995) 865–870.