

Zeitschrift: IABSE reports = Rapports AIPC = IVBH Berichte
Band: 58 (1989)

Artikel: Finding patterns in structural failures by machine learning
Autor: Stone, John / Blockley, David
DOI: <https://doi.org/10.5169/seals-44893>

Nutzungsbedingungen

Die ETH-Bibliothek ist die Anbieterin der digitalisierten Zeitschriften auf E-Periodica. Sie besitzt keine Urheberrechte an den Zeitschriften und ist nicht verantwortlich für deren Inhalte. Die Rechte liegen in der Regel bei den Herausgebern beziehungsweise den externen Rechteinhabern. Das Veröffentlichen von Bildern in Print- und Online-Publikationen sowie auf Social Media-Kanälen oder Webseiten ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. [Mehr erfahren](#)

Conditions d'utilisation

L'ETH Library est le fournisseur des revues numérisées. Elle ne détient aucun droit d'auteur sur les revues et n'est pas responsable de leur contenu. En règle générale, les droits sont détenus par les éditeurs ou les détenteurs de droits externes. La reproduction d'images dans des publications imprimées ou en ligne ainsi que sur des canaux de médias sociaux ou des sites web n'est autorisée qu'avec l'accord préalable des détenteurs des droits. [En savoir plus](#)

Terms of use

The ETH Library is the provider of the digitised journals. It does not own any copyrights to the journals and is not responsible for their content. The rights usually lie with the publishers or the external rights holders. Publishing images in print and online publications, as well as on social media channels or websites, is only permitted with the prior consent of the rights holders. [Find out more](#)

Download PDF: 12.02.2026

ETH-Bibliothek Zürich, E-Periodica, <https://www.e-periodica.ch>


Finding Patterns in Structural Failures by Machine Learning

Formulation de modèles de ruine des structures par apprentissage automatique

Erkenntnis von Mustern in Fehlleistungen durch «Machine Learning»

John STONE


Research Fellow
University of Bristol
Bristol, UK



John Stone is a Research Fellow in Civil Engineering at the University of Bristol. He is a Chartered Civil Engineer with ten years industrial experience in the design of engineering structures. His research interests include the application of artificial intelligence and machine learning to the problem of engineering safety.

David BLOCKLEY

Reader
University of Bristol
Bristol, UK



David Blockley is Reader in Civil Engineering at the University of Bristol. His research work has been concerned mainly with structural safety and reliability, with particular emphasis on the role of human error and the relationship with modern methods of risk and uncertainty analysis

SUMMARY

A new approach to the examination of significant features in structural failures is described. The method used is a development of the artificial intelligence (AI) technique of «machine learning» to extract sets of commonly occurring features from detailed reports of failures. By representing the information extracted hierarchically in the knowledge base of an expert system, advice can be obtained on the proneness to failure of a current project. The method is illustrated by comparison with a previous analysis of features in the failure of twenty-three engineering structures. A support logic measure of uncertainty is associated with each set of connected features.

RESUME

Cet article décrit une nouvelle approche de l'examen des caractéristiques significatives de la ruine des structures. La méthode utilisée est un développement de la technique d'intelligence artificielle (IA) nommée apprentissage automatique; cette technique met en évidence des caractéristiques communes décrites dans des rapports détaillés de ruine de structures. En présentant l'information ressortie hiérarchiquement de la base de connaissance d'un système expert, on peut obtenir des conseils à propos de la susceptibilité à la ruine d'un projet en cours. La méthode est illustrée par une comparaison, réalisée au cours d'un analyse précédente, de la ruine de vingt trois structures de génie civil.

ZUSAMMENFASSUNG

Ein neuer Weg zur Überprüfung von bedeutenden Merkmalen in strukturellen Fehlleistungen wird beschrieben. Die verwendete Methode, ist eine Entwicklung der künstlichen Intelligenz-Technik des «machine learning», um Gruppen der häufigsten Merkmale aus detaillierten Meldungen von Fehlleistungen zu extrahieren. Durch die Darstellung der hierarchisch extrahierten Information in der Wissensbasis eines Expertsystems, kann Mitteilung bezüglich der Neigung zu Fehlleistungen im laufenden Projekt erhalten werden. Die Methode wird mit einer früheren Analyse von Merkmalen in Fehlleistungen an dreiundzwanzig Ingenieurbauwerken verglichen.



1. INTRODUCTION

It is clearly important to study past failures and learn from them. The aim of the project reported in this paper has been outlined by Blockley [1] as the production of "...a knowledge based computer system which might be an aid in the management of the safety of a project". The central issues are a) the search for patterns in data concerning case histories, b) the need to handle uncertainty in open world problems and c) the use of the concept of a hierarchically structured knowledge base. In this paper we will concentrate on a) and only briefly outline b)[2] and c)[3,4].

2. DISCRIMINATION AND CONNECTIVITY

The search for patterns in data is commonly undertaken by the use of a variety of different types of cluster analysis [5]. These methods employ a number of different heuristics to determine 'groupings' of elements of data.

The methods of *discrimination* and *connectivity* described in this paper were developed initially by Norris, Pilsworth and Baldwin [6], who wished to investigate the relationship between medical symptoms and diseases from a number of patient case histories. They formulated two new methods of examining tabular numerical data in an attempt to overcome some of the theoretical and practical problems associated with the use of traditional clustering techniques.

The two methods are to be seen as complimentary approaches to the examination of relationships between features of objects and their classification (e.g. "symptoms" and "diseases"), but will be described here separately before presenting an example of their use in Section 3.

2.1 Discrimination

Discrimination entails the search for a *single* feature of an object which, by its presence or absence, gives evidence for the belief that an object belongs to one class rather than to another. It is therefore a *serial* approach.

The presence of discriminating features is common in engineering. For example, the range of feasible structural materials for the construction of a bridge might include reinforced concrete, steel and masonry, and the one chosen in a particular situation will generally depend upon a combination of factors. However, a requirement that the bridge should be movable for the passage of shipping would discriminate strongly in favour of the use of steel irrespective of the other competing factors. Low maintenance cost, on the other hand, might discriminate in favour of reinforced concrete whilst the matching of an adjoining masonry bridge might discriminate in favour of the use of masonry.

The initial stage in the discrimination analysis is to produce an incidence matrix I for each outcome where I_{ij} denotes the degree to which feature i was present in example or case j . For example, consider Table 1. This represents invented incidence data for two outcomes X and Y , each of which have three example cases. The examples may have one or more of the five features A - E . Note that in this example, all the I_{ij} values are either 0 or 1, denoting the certain absence or presence respectively of that feature. In a more general case, a multi-valued representation in the range [0,1] may be assigned, to represent a degree of belief.

A frequency distribution matrix F , as shown in Table 2, is then calculated, where f_{ij} denotes the proportion of those features i in outcome j , summed over all the cases.

		outcome					
		X			Y		
case j	feature i	1	2	3	1	2	3
A		1	1	1	0	0	0
B		1	1	0	0	1	0
C		1	1	0	0	0	1
D		0	0	1	1	1	1
E		0	0	0	0	1	1

Table 1 Incidence matrices I

feature i	outcome	
	X	Y
A	1	0
B	2/3	1/3
C	2/3	1/3
D	1/3	1
E	0	2/3

Table 2 Frequency matrix F

feature i	outcome	
	X	Y
A	1	0
B	1	0
C	1	0
D	2/3	1
E	0	1

Table 3 Discrimination matrix P

From the frequency distribution matrix the positive discrimination matrix P and the negative discrimination matrix N are calculated, using the definitions:

$$p_{ij} = \sum_{\substack{k \in D \\ k \neq j}} \left\{ x_{\text{ratio}}(f_{ij} / f_{ik}) \right\} / (C_D - 1) \quad \text{and} \quad n_{ij} = \sum_{\substack{k \in D \\ k \neq j}} \left\{ x_{\text{ratio}}(f_{ik} / f_{ij}) \right\} / (C_D - 1)$$

where $p_{ij}, n_{ij} \in [0,1]$ and C_D denotes the cardinality of the outcome set (i.e. the number of outcomes). The suffix D is used to denote a set of indices corresponding to the outcome set. The discrimination value is an accumulated measure of the degree to which the frequency of one feature is greater than that of all the other features for a given outcome. "ratio" is defined as a fuzzy set with membership characteristic function $x_{\text{ratio}} : R^+ \rightarrow [0,1]$ mapping the positive real numbers (i.e. f_{ij}/f_{ik} and f_{ik}/f_{ij}) onto the interval $[0,1]$. An example of such a fuzzy set is shown in Figure 1.




Fig. 1 Fuzzy set for discrimination analysis

Although the dashed line in Figure 1 represents a more general fuzzy set, the simplified solid line has been used for ease of computation. The resulting positive discrimination matrix for this example is given in Table 3. Note that if f_{ij} and f_{ik} are both equal to zero then $0/0$ is defined as equal to 0. Norris et al.[6] give a heuristic 'explanation' of the positive and negative discrimination measures which, when translated into the current terminology, argues that p_{ij} represents the accumulated belief that feature i is more indicative of outcome j than it is of any of the other outcome. This analysis therefore gives a method for assessing the significance of any single feature. The following section examines the importance of groups of features by using a connectivity analysis.



2.2 Connectivity

In contrast to the *serial* operation of the discrimination analysis, the connectivity algorithm adopts a *parallel* approach to the data. The method described here entails the search for *groups* of features which by their presence or absence give evidence for the belief that an object belongs to one class rather than to another. They are those features which have been found commonly to occur together and are associated with a given object classification. The algorithm is therefore a method for *pattern recognition*. Each outcome is considered in turn and a search is made for *groups* of features which commonly occur.

The connectivity analysis involves the calculation of both positive and negative connectivity matrices, with the negative analysis determining groups of features whose presence is indicative of the negation of a particular outcome. The starting point of the analysis is the incidence matrix I calculated during the discrimination phase. The positive connectivity analysis is applied to the incidence matrix directly whereas the negative connectivity analysis is applied to the *complement* of the incidence matrix.

If a and b are two feature vectors from an incidence matrix for a given outcome, then a connectivity measure, c_{ab} between a and b is defined as:

$$c_{ab} = \left\{ \sum_i (a_i \wedge b_i) \right\} / \left\{ \sum_i (a_i \vee b_i) \right\}$$

where \vee and \wedge denote maximum and minimum respectively and i ranges over the number of cases. The measure will be zero when a, b are disjoint and one when a, b are equivalent. Applying this algorithm to the incidence matrix in Table 1 for outcomes **X** and **Y** and the associated features, the positive connectivity matrices **C** of Tables 4 and 5 are obtained, with elements c_{ij} where i, j range over the feature names.

		feature				
		A	B	C	D	E
A	1	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{1}{3}$	0	
B	$\frac{2}{3}$	1	1	0	0	
C	$\frac{2}{3}$	1	1	0	0	
D	$\frac{1}{3}$	0	0	1	0	
E	0	0	0	0	1	

		feature				
		A	B	C	D	E
A	1	0	0	0	0	
B	0	1	0	$\frac{1}{3}$	$\frac{1}{2}$	
C	0	0	1	$\frac{1}{3}$	$\frac{1}{2}$	
D	0	$\frac{1}{3}$	$\frac{1}{3}$	1	$\frac{2}{3}$	
E	0	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{2}{3}$	1	

		feature				
		A	B	C	D	E
A	1	1	1	0	0	
B	1	1	1	0	0	
C	1	1	1	0	0	
D	0	0	0	1	0	
E	0	0	0	0	1	

Table 4 Positive connectivity matrix for outcome **X**

Table 5 Positive connectivity matrix for outcome **Y**

Table 6 Equivalence matrix from Table 4 for $\alpha = \frac{2}{3}$

Again in a more general case the incidence matrices will be multi-valued, with values in the range $[0,1]$. In this example, it can be seen that features **B** and **C** are the most strongly connected pair for outcome **X**, with $c_{BC} = c_{CB} = 1$, and **D** and **E** for outcome **Y** with $c_{DE} = c_{ED} = \frac{2}{3}$. This is intuitively to be expected from Table 1.

Having established a connectivity matrix **C** it is then possible to extract groups of connected features. This corresponds to finding paths in a graph [7]. A new relation can be derived by performing an α cut on the connectivity matrix. The new matrix contains values of 1 for those connectivities greater than or equal to α and zeros elsewhere. Warshall's algorithm [7] is then used to transform this symmetric matrix into a new connectivity matrix which can easily be partitioned. The partitions of the equivalence matrix now correspond to groups of features which are connected together at degree α . The value of α is set at various levels in the range $[0,1]$ and the resulting connected groups of tail-vertices examined. The equivalence matrix from Table 4 for $\alpha = \frac{2}{3}$ is as shown in Table 6. Two tables are produced for each outcome based on the positive

and negative connectivity analyses. Each table consists of sets of features corresponding to different values of the α cut in the range [0,1]. For example, the resulting table of connected groups of features from the positive connectivity matrix for outcome **X** in Table 4 is :-

$$\alpha = 1 : (\mathbf{B}, \mathbf{C}) \quad \alpha = \frac{2}{3} : (\mathbf{A}, \mathbf{B}, \mathbf{C}) \quad \alpha = \frac{1}{3} : (\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}) \quad \alpha = 0 : (\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E})$$

All the features form a single group at level $\alpha = 0$. At intermediate levels the features fall into separate groups, with the number of features in each group reducing as the connectivity level increases. Each higher level group may be thought of as being a representative set of features whose presence is evidence for the subsequent occurrence of the associated outcome (or evidence against it in the case of the negative table). In this example, the presence of group **(B,C)** is therefore strongly indicative of outcome **X**.

A pair of numbers in the range [0,1], known as a support pair [2], can be associated with a connected group at each connectivity level. These give lower and upper bounds on the evidential support for a proposition or event. The calculus is based upon an 'open world' representation of uncertainty in the sense that it is possible to represent propositions as true, false or unknown. The first number of the support pair, the necessary support, is given the value α from the connectivity analysis. The second number, the possible support, is always 1. Thus for outcome **X** a strong indicator is the group **(B,C)** with support pair (1,1). In support logic notation (a modified PROLOG rule) this is written **X :- (B,C) : [1,1]**. Other rules would be

$$\mathbf{X} : - (\mathbf{A}, \mathbf{B}, \mathbf{C}) : [\frac{2}{3}, 1] \quad \mathbf{X} : - (\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}) : [\frac{1}{3}, 1] \quad \mathbf{X} : - (\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}) : [0, 1]$$

The computer program implementing the connectivity method allows the step levels at which the connected groups are determined (in the range $0 \leq \alpha \leq 1$) to be chosen by the user. This enables the structure of the groups at different connectivity levels to be examined as appropriate to the application. The values $0, \frac{1}{3}, \frac{2}{3}, 1$ have been selected for this example since the small number of feature groupings are portrayed adequately. A more complex example in Section 3 illustrates a closer division at increments of 0.1.

3. FAILURE ANALYSIS BY SIMPLE SUMMATION

A detailed account of a simple analysis of structural failures has been given previously by Blockley [8]. This paper describes the application of the connectivity and discrimination analyses to Blockley's data, which were assessments of the relative truth (or dependability) and importance of a number of statements concerning well documented failures. The original investigation examined twenty four statements about twenty three failures, ranging from the Tay Bridge collapse of 1879 to the loss of the oil drilling barge Trans Ocean 3 in the North Sea in 1974. Typical statements were "the structure is not sensitive to random hazards", "the designers are adequately experienced in this type of work" and "the contractual arrangements are perfectly normal". The assessments, although made with engineering judgement and experience, were entirely subjective and personal, made with the benefit of hindsight, and a different investigator may have chosen quite different values. They do, however, represent a useful basis for analysis since they are a consistent set of interpretations carried out with a common method and purpose.

The assessments were made on the basis of five categories for both truth and importance. The level of confidence in the *truth* of a statement was graded between 1 (very high) and 5 (very low), and its *importance* between A (very low) and E (very high). Numerical values were assigned to each of the assessment ratings 1-5 and A-E on the following scale: 1 and A = 0.2, 2 and B = 0.4, 3 and C = 0.6, 4 and D = 0.8, 5 and E = 1.0. An overall 'combined rating' was calculated as the product of the two individual values and lying in the interval [0,1]. Thus a rating of 4 for truth and C for importance yields a combined figure of 0.48. The new combined assessments are therefore



represented in the form of a fuzzy incidence matrix as described in Section 3.1. A simple analysis was made [8] by summing the combined ratings over all twenty three failures. Table 7 is a short extract from the results.

Order	Statement	Brief description	sum
1	5a	design error	15.48
2	6c	construction error	11.88
3	6e	contractor's staff	11.76
4	5d	designer's site staff	11.68
5	2b	R & D information	10.88

Table 7 Simple summation of accident statement parameters

These values have been interpreted elsewhere [8], but it is important to emphasise here that the sample of failures from which they were derived is not random and includes only failures important enough to merit individual reports of inquiry. The scores are not precise numerical quantities but only relative indications of the importance of the statements.

4. FAILURE ANALYSIS BY DISCRIMINATION AND CONNECTIVITY

The following example illustrates the application of the discrimination and connectivity methods to the analysis of the failure data given in Section 3. Note that all the cases selected refer to failures, so at one level there is therefore only one outcome to be considered. This means that a discrimination analysis cannot be performed, since it by definition determines an ability to discriminate *between* outcomes. However, it is apparent that the cases may also be classified by other means, such as mode of failure (e.g. fatigue, overstress etc.), structural form (e.g. bridge, oil-platform) or other criteria as desired. With more than one outcome, a discrimination analysis may be made. For example, if the cases are partitioned into 'bridges' and 'others' then the positive discrimination matrix shown in Table 8 is obtained.

state- ment	bridges	others	state- ment	bridges	others	state- ment	bridges	others
1a	0.00	0.23	4c	0.00	0.22	6a	0.39	0.00
1b	1.00	0.00	4d	0.13	0.00	6b	0.40	0.00
2a	0.00	0.18	4e	0.09	0.00	6c	0.00	0.11
2b	0.15	0.00	5a	0.00	0.04	6d	0.44	0.00
3a	0.04	0.00	5b	0.00	0.00	6e	0.14	0.00
3b	0.00	1.00	5c	0.83	0.00	7a	0.00	0.00
4a	0.02	0.00	5d	0.10	0.00	7b	0.54	0.00
4b	0.00	0.48	5e	0.36	0.00	8	0.00	0.36

Table 8 Positive discrimination values

The reader is referred to the original analysis [8] for the full list of the meaning of each of the statements 1a - 8. It is apparent that statement 1b (strength variability) discriminates strongly in favour of a bridge failure, and statement 3b (sensitivity to random hazards) in favour of 'other' failure. From the original data [8] it can be seen that the only occasions on which statement 1b was assessed as being of other than minimal significance both related to bridge failures (Tay and Quebec 2). It never appeared as a significant factor in any of the 'other' failures, and therefore discriminates in favour of bridge failures. A similar argument applies to statement 3b, which only occurred as a significant factor once, relating to an 'other' failure (Ronan Point).

From the discrimination analysis the frequency of occurrence of a feature for an outcome is obtained *relative to all the other possible outcomes*. Thus even a single occurrence can be significant if it occurs only for one outcome and never for any of the others, as seen above with statement 3b. A connectivity analysis carried out on the same data gives the following groups of positively connected statements:

$\alpha = 0.5$:	(5d, 6e)
$\alpha = 0.4$:	(5d, 6c, 6e)
$\alpha = 0.3$:	(5a, 5d, 6c, 6e)
$\alpha = 0.2$:	(2a, 2b, 5a, 5d, 6c, 6e)
$\alpha = 0.1$:	(1a, 2a, 2b, 4a, 4b, 4d, 4e, 5a, 5d, 5e, 6c, 6e, 7a, 7b, 8)

Note that no statements are connected more strongly than at a connectivity level of 0.5. This reflects the diverse nature of the practical problem, and is in contrast to the artificially chosen example of Section 3.2 where some features were connected at $\alpha = 1.0$. It is interesting to compare the above connected groups with the results from the previous analysis shown in Table 7. The most strongly connected pair of statements from the connectivity analysis, 5d and 6e, occur as the fourth and third most frequent statements respectively in Table 7. Two connectivity intervals lower, at $\alpha = 0.3$, all the top four of the previous results - 5a, 6c, 6e and 5d - are now found to be connected in one group.

The *meaning* of the connected groups may be considered as a new entity rather than in terms of the individual features. For example, the two most strongly connected statements (5d and 6e) each relate to a specific aspect of site control staff. "Site control staff" may therefore be thought of as a *higher level*, or more general, description encompassing both of the statements. Similarly, the second level group (5d, 6c and 6e) adds "construction error" to the top level group, and might therefore be thought of as "site procedure". Grouping together concepts in this fashion leads to the possibility of constructing a hierarchical knowledge base, consisting of progressively more general concepts at higher levels, and grounded in specific concepts relating to individual case histories at the bottom.

5. HIERARCHICAL KNOWLEDGE BASE

The analysis based upon connectivity, discrimination and the grouping together of related concepts can be extended [9] to develop a hierarchically structured knowledge base of case histories of failures. The case histories are 'captured' by the use of event sequence diagrams (ESDs) [10]. These diagrams show the temporal order and relationship of events leading up to a particular outcome. For example, Figure 2 shows a hierarchy of three ESD representations of the same 'story'. The lowest level (Level 1) is the most detailed, corresponding to specific detailed concepts from a case history. Level 2 is an intermediate representation, in which a number of the bottom level concepts have been 'merged' to form new, broader concepts as noted above. The most general representation, Level 3, is obtained by further merging of concepts by repeated application of the connectivity analysis.

The advantage of representing knowledge in a hierarchical form is realised when it is wished to query the knowledge base. The user is able to pursue a query about a concept to an appropriate level of detail.

Figure 3 is an outline of the structure of our proposed development of a knowledge-based system (KBS) to fulfil the objective stated in the Introduction. The upper section of the diagram, concerned with building the knowledge base, has been implemented in C on an IBM PC/AT. The lower section illustrates the proposed future development and use of the system. 'FRISP' is 'Fuzzy Relational Inference with SuPport logic' [11], a PROLOG program which allows a user to query a



knowledge base and to receive information with the associated uncertainty expressed in the form of a support pair. The structure of the KBS includes two 'learning' loops. The first is in the upper, 'building' section, where knowledge from case histories is accumulated and merged into the hierarchical knowledge base. The second joins the 'building' and 'user' sections, and commences when the representation of the 'world' embodied in the knowledge base and the FRISP responses become unsatisfactory. This situation can only be remedied by the user providing the 'building' phase of the system with more case histories, which may result in the formation of new concepts in the knowledge base following a new connectivity analysis.




Figure 2 Hierarchy of Event Sequence Diagrams




Figure 3 Outline of structural safety KBS

6. CONCLUSIONS

The application of a method of machine learning based upon the techniques of discrimination and connectivity to the assessment of structural safety has been described. The following points are significant :-

- 1) Many engineering failures contain common features. Attempts to learn from failures may therefore be based upon pattern recognition techniques.
- 2) The detection of common patterns of features may consider either single features (discrimination) or groups of features (connectivity).
- 3) The support logic calculus allows an appropriate 'open world' representation of uncertainty, in which propositions are either true, false or unknown.
- 4) The use of the connectivity analysis allows the meaning of groups of features to be merged in a hierarchical form suitable for inclusion in a structural safety knowledge base.



ACKNOWLEDGEMENTS

The work reported in this Paper was supported by a grant from the United Kingdom Science and Engineering Research Council.

The authors wish to thank Rosie Greenhill and Mrs V.M. Jenkins for their help with the translations.

REFERENCES

1. BLOCKLEY, D.I. An A.I. tool in structural safety control. In A.S. Nowak (Ed.), *Modelling human error in structural design and construction*. New York: A.S.C.E., 1986.
2. BALDWIN, J.F. Evidential support logic programming. *Fuzzy Sets and Systems*. **24**, 1987, pp 1-26.
3. STONE, J.R., BLOCKLEY, D.I. and PILSWORTH, B.W. Towards machine learning from case histories. *Civil Engineering Systems* (in press).
4. STONE, J.R. *Expert system learning from structural failures*. Ph.D. thesis, University of Bristol, 1989.
5. EVERITT, B. *Cluster analysis* (2nd ed.). London: Heinemann, 1980.
6. NORRIS, D., PILSWORTH, B.W. and BALDWIN, J.F. Medical diagnosis from patient records - a method using fuzzy discrimination and connectivity analyses. *Fuzzy Sets and Systems*, **23**, 1987, pp 73-87.
7. LEVY, L.S. *Discrete structures of computer science*. New York: John Wiley, 1980.
8. BLOCKLEY, D.I. *The nature of structural safety and design*. Chichester: Ellis Horwood, 1980.
9. STONE, J.R. and BLOCKLEY, D.I. Structural reliability through learning from case histories. *Proceedings of Fifth International Conference on Structural Safety and Reliability* (ICOSSAR '89), San Francisco, 7-11 August. A.S.C.E.
10. TURNER, B.A. *Man-made disasters*. London: Wykeham, 1978.
11. BLOCKLEY, D.I. and HENDERSON, J.R. Knowledge base for risk and cost benefit analysis of limestone mines in the West Midlands. *Proc. Instn. Civ. Engrs.*, Part 1, **84**, June, 1988, pp 539-564.