A CASE-BASED REASONING APPROACH FOR SELECTING RISK MANAGEMENT TECHNIQUES

Doug Forbes[1], Simon D. Smith[2] and Malcolm Horner[3]

[1,3]Construction Management Research Unit, Division of Civil Engineering, University of Dundee, Dundee, DD1 4HN, UK
[2] Institute for Infrastructure and the Environment, School of Engineering & Electronics, University of Edinburgh, Kings Buildings, Edinburgh, EH9 3JL, UK

The management of risk in the built environment is critical and there is a wide range of techniques available to deal with the task. However, research has shown that only a small number of techniques are used by practitioners (Akintoye and MacLeod 1997, Wood and Ellis 2003). One reason cited for this is a lack of knowledge of the circumstances in which they can be used. The aim of the research in this paper is to produce a case-based reasoning (CBR) tool for decision support in selecting appropriate techniques for built environment problems. The tool uses a case-base developed from historic problems in the literature. A problem framework is used to characterize problems as 1) External or internal – described by the PESTLE Model, 2) The Risk Owner and 3) The Project Phase. A second stage defines the data used in a problem as fuzzy, incomplete or random. The methodology of CBR is heuristic and as such this work has investigated the effect of differing the retrieval mechanism, the inclusion of weights and the threshold value. The results demonstrated a tool which during validation predicted the correct technique up to 93% of the time; additionally it was seen that more complex CBR methods did not result in more accurate prediction rates. Overall the research has produced a simple tool to select appropriate risk management techniques and demonstrated the applicability of CBR to the problem – highlighting that the simplest methodology has proved the most effective.

Keywords: risk techniques, decision support, case-based reasoning.

INTRODUCTION

The construction industry needs to manage risk (Flanagan and Norman 1993), and there is a wide range of techniques available to do so. The work presented in this paper creates a decision support tool, using case-based reasoning (CBR), to select appropriate risk techniques for the built environment. The purpose of this tool is to facilitate the selection of techniques from the broad range which is available for application in the built environment.

Risk Management

The research is concerned with techniques applied in the first three stages of the risk management cycle - identification, estimation (analysis stage I) and evaluation (analysis stage II) (Perry and Hayes 1985, Flanagan and Norman 1993, Baker *et al.* 1999). The response and monitoring stages are not considered here, as specific techniques are not applied.

[1] d.r.forbes@dundee.ac.uk

Forbes, D, Smith, S D and Horner, M (2007) A case-based reasoning approach for selecting risk management techniques. *In:* Boyd, D (Ed) *Procs 23rd Annual ARCOM Conference*, 3-5 September 2007, Belfast, UK, Association of Researchers in Construction Management, 735-744

The literature covering construction management risk techniques has shown a reliance in practice on a small number of techniques (Akintoye and MacLeod 1997, Bajaj *et al.* 1997). This is in contrast to the 36 identified in the literature, with most practitioners using simple methods and avoiding 'complex' ones. Most complex techniques (eg. stochastic dominance, influence diagrams, artificial intelligence), were rarely even considered by the practitioners (Wood and Ellis 2003). Many felt that they did not know when to apply a technique or if they were appropriate to the built environment (Akintoye and MacLeod 1997). This is perplexing as the literature contains many examples of the 'complex' techniques being applied to construction problems, demonstrating their applicability. The need for a decision support tool for selecting risk management techniques has been identified in the literature (Dikmen *et al.* 2004, Wang *et al.* 2004). It is proposed, through this work, to develop a CBR tool which will be based on examples from the literature of the application of techniques and suggest techniques which could be applied for similar sets of problems characteristics.

Case-based Reasoning

Case-based reasoning applies what has taken place in the past and infers its application to a new situation (Kolodoner 1993). The output solution produced is based on the similarity of the features of the new case to cases in the case-base. Thus, if a new case has all the features matching a historic case a similarity of 100% will be obtained. Case-based reasoning is not an artificial intelligence technology, although often referred to as such (Watson 1999). Instead, it is a methodology which can be applied to solve knowledge based problems. As such the stages which exist in the development of a CBR model are dependent on the problem to which they are applied. The stages central to the application of CBR are: definition of the domain; representation of the cases; indexing and storage; retrieval of the cases and adaptation as required (Karim and Adeli 2003). The principles of CBR in being able to solve new problems based on historic cases make it suited to selecting appropriate risk management techniques.

BUILDING THE CASE-BASE

The Problem Breakdown Structure

The case-base was created from examples in the literature of technique applications. To assess the characteristics of each example for a given technique a standard breakdown structure was devised. This structure has two parts. The first (A) examines the nature of the problem and the second (B) the nature of the data used. Stage A initially assesses a problem as being either external or internal to the organization for whom the risk is being assessed. This is then split into Political; Economic; Socio-cultural; Technological; Legal and Environmental (PESTLE) elements, which have been previously applied to risk management (HM Treasury 2004). Stage A then records the risk owner – defined as being contractor; client; consultant/designer; financier; facilities management organization or government. The list of risk owners was created from those identified in risk management problems. Others could be added if necessary. Finally, the project phases are defined as (1) Inception/Feasibility; (2) Design; (3) Construction; (4) Commissioning; (5) Operation; and (6) Decommissioning. These phases are based on the RIBA Plan of Work (2000) and whole life costing (Marenjak 2004).

In contrast to the problem characteristics in Stage A, Stage B took account of the characteristics of the data used in the given example. The FIR Model (Fuzziness, Incompleteness and Randomness) was used for the assessment (Blockley and Godfrey 2000). Fuzziness is the imprecision of definition (Blockley 1995). For the purposes of defining the nature of data in a risk management problem, fuzziness is present if that which is being assessed is imprecise. For example fuzziness may occur in terms such as high, medium or low; or large and small costs. The incompleteness aspect of a model is concerned with that which is not known (Blockley and Godfrey 2000). Using this basis all risk management models are attempting to model incompleteness. To overcome this, the exact nature of the incompleteness was refined to relate solely to the data that was applied in the model. It was found that this provided a better representation of the incompleteness of the data. Finally, randomness is the lack of a specific pattern in the information (Blockley 1995). This is the uncertainty defined by probability and statistics (Blockley and Godfrey 2000), and has been assigned for each of the techniques by assessing data with no specific pattern.

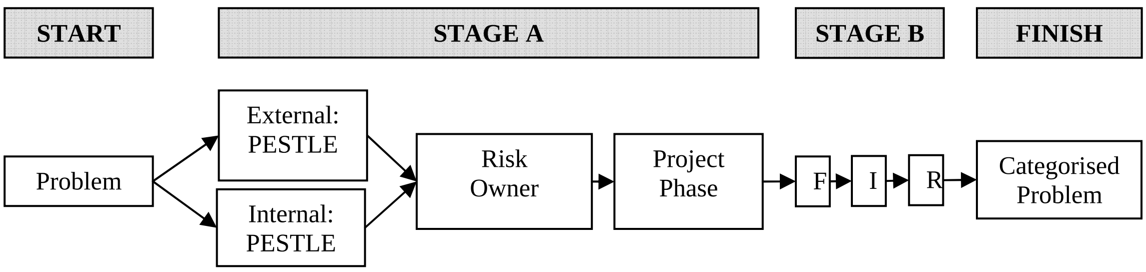

Figure 1: The Problem Breakdown Structure

The final breakdown structure is given in Figure 1. This shows the two stages, comprising the three elements of A and the F, I & R elements of B. These define the six 'input' features which will be used to define each case.

Defining the Indices

While the six 'input' features are defined above, the relevant risk management technique is the 'output' feature. There were 52 techniques identified from the literature over the three stages. These have been grouped into 23 categories as shown in Table 1, to overcome the issue of such a large number.

Table 1: Risk Management Categories Defined for the Matrix

Identification		**Estimation**		**Evaluation**	
1	Artificial Intelligence	8	Artificial Intelligence	17	Artificial Intelligence
2	Decomposition	9	Decomposition	18	Decomposition
3	Experiential	10	Experiential		
4	Failure ID	11	Failure ID		
		12	Other	19	Other
		13	Probabilistic	20	Probabilistic
5	Sensitivity analysis	14	Sensitivity analysis	21	Sensitivity analysis
6	Support Systems	15	Support Systems	22	Support Systems
7	Trees	16	Trees	23	Trees

Each of the 23 categories have been assigned to the appropriate three stages of the risk management cycle. An index of 1-23 consecutive numbers have been applied to each. This allows the CBR model to select the appropriate technique at a given stage. Indices for the 'input' features did not need to be grouped, because of the small number of possible permutations. Thus, the index for each input feature was directly related to the range of values in each. For example, the 12 elements of the external and internal PESTLE assessment were assigned 1-12. These values are nominal and, for instance, 12 in this feature is as different from 11 as from 1.

Analysis of Techniques

The problems in the literature were assessed against each of the six ‘input’ features, each techniques used was assigned one of the 23 categories. The process examined 179 examples from built environment problems. During this process two issues were identified. The first was that one technique may be applied, for one risk owner or stage, to more than one problem characteristic. In order to overcome this the permutations of each of the problems characteristics (PESTLE; Owner; Stage; FIR) for a given example of a technique were created to maintain 6 features. This also had the advantage of creating a larger case-base. The second issue is that not all of the examples refer to an actual application (either real or hypothetical). Some merely suggested situations for which a technique could be applied. These were recorded separately as it was not considered they should be given the same weight as if an example were provided. The assessment of the 179 examples resulted in 6177 cases in the case-base after the permutations of each had been derived.

DESIGN OF THE MODELS

In order to develop the most appropriate models and to ensure that the case-base is representative of the nature of the risk management techniques, modifications to the contents of the case-base were investigated. The first of these addressed the ‘suggested’ applications. For the development of the model there were three options:

1. Re-classify the ‘suggested’ cases be equivalent to ‘actual’ cases;
2. Remove the ‘suggested’ cases from the case-base and deal solely with the ‘actual’ cases;
3. Design the model in such a way that the weighting associated with the ‘suggested’ cases is not as high as ‘actual’ values.

Options 1 and 2 are concerned with the data in the case-base and are therefore the most easily investigated. The third involves modifying the CBR system, and has the potential problem in defining the level of the weights which should be applied. The first two were therefore investigated, leaving a decision to be made on the third if adequate results were not obtained.

The second issue was assessing the impact of each of the three stages of the risk management cycle (identification; estimation; evaluation). The index that was applied to this should provide the reasoner with the ability to produce the appropriate technique. However, initial trials suggested that the wrong stage may be selected if some problem characteristics exist for more than one stage. To overcome this, the case-base was split into each of the three stages. A case-base containing all three stages combined was retained to allow a comparison.

As a result of these investigations 8 models were proposed: four to cover all of the three stages separated, plus the combined; which were then repeated for the two options of dealing with ‘suggested’ values.

CASE RETRIEVAL

The heuristic nature of the CBR methodology is a drawback in its application. The research carried out for this work investigated four elements: 1) the feature similarity; 2) weighting mechanism; 3) selection method; 4) threshold value.

Feature Counting

To assess the similarity of a new case to the cases in the case-base, the similarity has to be measured. One of the most frequently applied methods is 'nearest neighbour' (Watson 1997, Chiu 2002). This method implies that each case is a dimension in a domain space and that the similarity of each of the cases in the case-base to the new case is shown in Equation 1. In this equation, *f* – the similarity of the feature *i*– can be modified depending on how the similarity is measured (Watson 1997).

$$Similarity(\%) = \frac{\sum_{i=1}^{n} w_i f_i}{n} \times 100 \qquad \text{Equation 1}$$

Where *n* is the number of features, *w* is a weight applied to each feature.

For the purposes of this research the value of *f* is 1 if the two input features match and zero otherwise. No other measures of similarity have been used because of difficulties in measuring between two nominal values.

Weighting Mechanism

It is clear that some features may be more important in defining the output than others. Previous researchers have investigated ways by which weights can be applied to CBR. This research applied the ID3 Algorithm (Watson 1997, Arditi and Tokdemir 1999, Graham 2005). This algorithm was designed to construct complex decision trees and is based on the principle of information theory (Quinlan 1986). ID3 is able to define, at each node of the tree, which attribute should be applied – ie, which attribute is the best classifier (Mitchell 1997). The relative importance of each of these nodes is then used to define the feature weights. The risk management selection tool model has investigated the application of ID3 weights to the six features and compared this to an assumed equal weighting of all features.

Case Selection

After each of the cases in the case-base has been assigned a similarity score a method is required to select the most appropriate case. This area is not covered in detail in the literature as most researchers have selected the case with the highest similarity score. A study by Arditi and Tokdemir (1999) developed three final prediction methods:

Method I: *The outcome of the retrieved case with the highest similarity score is selected (MI)*

Method II: *The outcome that appears most frequently is selected from among the outcomes of the first 10 cases that had a similarity score above the threshold (MII)*

Method III: *The outcome of the cases that received the highest average similarity score is selected from among the outcomes of the first five cases that had a similarity score above the threshold (MIII)*

Arditi and Tokdemir (1999) identified that Method I produced the best results and that Method II poorer results. Method I is the simplest to implement, is less applicable where there are a small number of features – 6 in this case. This increases the chance of several equally high scores. For this reason a fourth method (MIV) was developed which uses the most commonly occurring case from the multiple solutions of Method I, II and III. Due to the sorting of the data as nominal values in the CBR risk tool the first 10 and first 5 used in Methods II and II respectively have been modified to include all the values above the threshold.

All four methods were investigated in the course of this research. During the testing and validation stages, because more than one case may be selected by each of the above methods, if any of these matched the test or validation cases it was considered a full match.

Threshold Value

The final element of the methodology which was varied was the threshold value. Commonly this is set at 75 % (Arditi and Tokdemir 1999). It was noted that for the six features in the model at least five would have to match to score the 75%. There is a strong possibility, therefore, that no cases are selected. To assess variations in threshold level, it was also trialled at 66% and 0%.

TESTING AND VALIDATION

Literature relating to validation of CBR models is limited (O'Leary 1993, Gonzalez *et al.* 1998, Graham and Smith 2004). The testing and validation method used for the risk management selection tool has been adapted from neural networks and follows similar principles to that of O’Leary (1993). The method is the ‘train and test’ validation process (Twomey and Smith 1997, Shi 1999). This splits the data into three: a train, test and validation set. The network is built using the train set, and the testing set gives an initial indication of performance. The validation set is independent and gives an indication of the predictive capability of the model. However, the testing data allows an initial assessment to be made of the suitability of models, such as the split into the three stages of the risk management cycle.

The validation set was developed from further literature example to ensure independence. There were 42 examples, 20 of which contained ‘suggested’ rather than actual examples. Following the same permutation process as for the case-base, 1226 cases were included when the ‘suggested’ cases were considered to be ‘actual’ and 476 when ‘suggested’ examples were excluded. These validation sets were then applied to the combined case-bases of training and testing data.

RESULTS

Four basic models were built using different case-bases. The prediction rates for the valuation sets are shown in Table 2 covering the full case-base (A); identification (B); estimation (C) and evaluation (D) techniques. The procedure was repeated for the inclusion of ‘suggested’ values and the exclusion of these. This comprised the eight models. These was repeated for three levels of threshold value and included all four retrieval mechanisms and the different feature weights.

The results from this validation stage demonstrated the following:

- Statistically better results were obtained at 95% confidence when the data was split into the three stages of the risk management cycle;

- Weighting the features produced statistically poorer or the same results (at 95% confidence) in 23 out of 96 cases;
- Statistically better or the same results were obtained using Method I over other retrieval mechanisms;
- Methods III and IV resulted in poorer outcomes than Method I and II for all the models. Method III failed to predict any correct techniques in 16 out for 48 cases – and never predicted more than 43% correctly;
- The results achieved were better for the data set with the suggested values removed than when the suggested values were assumed to be equal to the actual;
- The best set of results was achieved using the data split into the three stages with the suggested values omitted. This was done using Method I retrieval. The non-weighted values were statistically better at 95% confidence for the identification techniques and the same for estimation and evaluation.

Table 2: Validation Prediction Rates for CBR Tool (all values in %)

	Thres-hold		Non-Weighted MI	Non-Weighted MII	Non-Weighted MIII	Non-Weighted MIV	Weighted MI	Weighted MII	Weighted MIII	Weighted MIV
Suggested Values Included	75%	A	42.4	21.5	5.9	21.4	38.9	27.8	0.0	25.9
		B	72.8	63.6	15.8	60.1	44.3	88.2	3.9	43.4
		C	60.9	38.8	27.9	38.9	55.8	54.2	32.6	53.7
		D	53.9	46.9	20.2	46.6	61.2	59.6	0.0	59.3
	66%	A	42.4	17.9	1.7	17.9	38.9	25.4	0.0	23.5
		B	72.8	88.2	3.9	71.1	44.3	88.2	3.9	43.4
		C	60.9	46.7	25.4	49.1	55.8	46.1	22.9	54.0
		D	53.9	45.5	5.9	45.2	61.2	59.6	0.0	59.3
	0%	A	42.4	16.4	0.0	14.4	38.9	16.4	0.0	14.2
		B	72.8	88.2	7.5	77.2	44.3	65.4	3.9	43.4
		C	60.9	18.7	22.9	20.2	55.8	18.7	22.9	22.1
		D	53.9	0.0	0.0	0.0	61.2	0.0	0.0	0.0
Suggested Values Excluded	75%	A	62.8	33.6	13.7	33.2	53.4	37.8	0.0	32.6
		B	87.1	54.0	21.0	54.8	79.0	64.5	7.3	54.8
		C	93.0	68.1	43.2	68.1	91.4	90.3	39.5	90.3
		D	92.8	75.4	34.7	74.9	93.4	91.0	0.0	90.4
	66%	A	62.8	34.5	6.3	33.0	53.4	33.6	0.0	30.0
		B	87.1	64.5	7.3	59.7	79.0	50.8	7.3	36.3
		C	93.0	90.3	33.5	90.3	91.4	90.3	9.7	90.3
		D	92.8	91.0	18.0	90.4	93.4	91.0	0.0	90.4
	0%	A	62.8	35.1	0.0	31.7	53.4	35.1	0.0	33.4
		B	87.1	92.7	7.3	83.1	79.0	92.7	7.3	75.0
		C	93.5	90.3	9.7	90.3	91.4	90.3	9.7	90.3
		D	92.8	91.0	0.0	89.8	93.4	91.0	0.0	90.4

DISCUSSION

The validation results show that the most suitable models are developed using the case-base which excludes the suggested values and splits the data into three stages. The best outcomes were generally obtained at 66% and 0% thresholds.

The inclusion of the suggested examples was shown to have a significant impact on the outputs. It is interesting to note that the complete exclusion of these values resulted in better results. There are two possible reasons for this, firstly it may be that these examples are suggested as a panacea in the literature, but are not. Secondly, in assessing the actual example more details are given simply because they are examples – thus the exact nature of the category is easier to assess. This results indicate a model which can be used without investigating the suggested values further. Consequently there can be no real advantaged in investigating the third option outlined - to re-design the reasoner taking account of the relevant merit of the two types of examples.

The split of the data into the three stages also has an impact on the results. It was originally intended that this would be dealt with by the '1-23' identifier tag. However, this did not safeguard against selecting techniques for another stage. The split was therefore required to overcome this, however this will not be an issue when using the tool for decision support as the required stage will be known.

When the threshold level was lowered from 75 to 66%, the prediction rate increased. A further decrease to 0% did not make a difference to 24/32 of the cases, increased it in 5 and decreased it in 4. This demonstrated that the model is subject to variations in the threshold value and would suggest that a decrease in level results in a generally higher prediction rate. However, this value is only critical for Methods II-IV.

From the four retrieval mechanisms, Method I was consistently the best performer for the non-weighted, and predominantly for the weighted. During the validation process, if the tool returned more than one technique with an equally high score matching the validation case, a match was deemed to exist. These equally high scores for different techniques mean that more than one technique has been applied to the problem in the past. By examining the range of techniques that are returned, the decision maker can choose a technique to apply to the current situation. It is interesting to note that Method I is the simplest of all four, and leaves a degree of responsibility with the decision maker. The other methods use more complex solutions (using mode and mean) and it can be seen that they do not always lead to the selection of the correct technique. It is possible that these methods may need a further, or combination of, methods to be successful. However, the decision is likely to be based on other factors such as familiarity with the techniques which can be made when a full list of suitable techniques is presented.

This highlights the limitations of the case-base. The case-base has been developed from 179 examples in the literature. As the case-base grows the outputs from Methods II-IV may differ. Method I, in contrast, will grow because it incorporates all the new cases to the output. This stresses the important point that just because a technique is not returned as being appropriate by Method I does not mean it necessarily cannot be used. It simply means it is not in the case-base.

A final element for discussion is the weighting mechanism. It has been shown through the results obtained that the application of the ID3 weights to the case-base tends to provide poorer results. It is suggested that a reason for this may be that the weights cannot be applied evenly to all of the features in every case, and so the approximation of equal weighting for every case provides better results.

Further work

Following the production of a tool which produces good results, further work has been identified relating to deepening the knowledge of CBR. The first area will be to assess the impact of the retrieval mechanism and the threshold. This should be coupled with the impact of the weights. A further area of research needed will be to measure similarity of features which have nominal, rather than ordinal or interval, measures.

This research has focussed on applying feature counting, but other methods could be adapted and investigated from clustering techniques to CBR.

CONCLUSIONS

In conclusion, a CBR model has been produced which can predict appropriate risk management techniques for use in the built environment. The tool makes an assessment based on the frequency of occurrence from historic cases, and allows decision makers to choose an appropriate technique. However, it is still necessary for the user to gain knowledge of the techniques. The problem breakdown is constrained by the use of the standard breakdown; but this has the advantage of being simple while still being comprehensive. The validation process has shown it to predict the correct solutions in over 80% of cases.

The CBR tool development has confirmed the heuristic nature of the methodology, and highlighted the need to investigate many avenues when applying CBR. It is particularly interesting to note that the methodology which has proved the most successful is that which keeps the process simple – no weights and selecting the highest similarity scores. However, despite this, the work has highlighted that the field of CBR would benefit from gaining a greater understanding of the way in which a match against nominal values in the case-base can be evaluated.

REFERENCES

Akintoye, A S and MacLeod, M J (1997) Risk Analysis and Management in Construction. *International Journal of Project Management*, **15**(1), 31-8.

Arditi, D and Tokdemir, O B (1999) Using Case-Based Reasoning to Predict the Outcome of Construction Litigation. *Computer-Aided Civil and Infrastructure Engineering*, **14**, 385-93.

Bajaj, D, Oluwoye, J and Lenard, D (1997) An analysis of contractors' approaches to risk identification in New South Wales, Australia. *Construction Management and Economics*, **15**(4), 363-9.

Baker, S, Ponniah, D and Smith, S (1999) Risk response techniques employed for major projects. *Construction Management and Economics*, **17**(2), 205-13.

Blockley, D (1995) Process Re-engineering for Safety. In: James, M (Ed.), *Risk management in civil, mechanical and structural engineering*, 22 February, London, 51-66.

Blockley, D and Godfrey, P (2000) *Doing It Differently: Systems for Rethinking Construction*. London: Thomas Telford Ltd.

Chiu, C (2002) A case-based customer classification approach for direct marketing. *Expert Systems with Applications*, **22**(2), 163.

Dikmen, I, Birgonul, M T and Arikan, A E (2004) A Critical Review of Risk Management Support Tools. *In: 20th Annual Conference of Association of Researchers in Construction Management (ARCOM)*, September 1-3, Heriot Watt University, Edinburgh, Vol. 1, 1145-54.

Flanagan, R and Norman, G (1993) *Risk Management and Construction*. Oxford: Blackwell.

Gonzalez, A J, Xu, L and Gupta, U M (1998) Validation Techniques for Case-Based Reasoning Systems. *In: IEEE Transactions on Systems, Man and Cybernetics-Part A: Systems and Humans*, 1998, Vol. 28, 465-77.

Graham, D (2005) *Stochastic modelling and analysis of construction processes*, PhD Thesis, University of Edinburgh.

Graham, D and Smith, S D (2004) Estimating the Productivity of Cyclic Construction Operations Using Case-based reasoning. *Advanced Engineering Informatics*, **18**, 17-28.

HM Treasury (2004) *The Orange Book. Management of Risk-Principles and Concepts*. Norwich: HMSO.

Karim, A and Adeli, H (2003) CBR Model for Freeway Work Zone Traffic Management. *Journal of Transportation Engineering*, **129**(2), 134-45.

Kolodoner, J (1993) Case-Based Reasoning. San Mateo, California, USA: Morgan Kaufmann.

Marenjak, S (2004) *A generic approach to minimizing WLC in the building industry*, PhD Thesis, University of Dundee.

Mitchell, T M (1997) *Machine Learning*. New York: McGraw-Hill.

O'Leary, D E (1993) Verification and Validation of Case-Based Systems. *Expert Systems with Applications*, **6**, 57-66.

Perry, J G and Hayes, R W (1985) Risk and its management in construction projects. *Proceedings of the Institution of Civil Engineers*, **78**(1), 499-521.

Quinlan, J R (1986) Induction of Decision Trees. *Machine Learning*, **1**, 81-106.

RIBA (2000) *Plan of Work*. London: RIBA Enterprises.

Shi, J J (1999) A neural network based system for predicting earthmoving production. *Construction Management and Economics*, **17**(4), 463-71.

Twomey, J M and Smith, A E (1997) Validation and Verification. *In*: Kartam, N, Flood, I and Garrett, J H (Eds.), *Artificial Neural Networks for Civil Engineers: Fundamentals and Applications*. New York: ASCE.

Wang, J, Sii, H S, Yang, J B, Pillay, A, Yu, D, Liu, J, Maistralis, E and Saajedi, A (2004) Use of Advances in Technology for Maritime Risk Assessment. *Risk Analysis*, **24**(4), 1041-63.

Watson, I (1997) *Applying Case-based Reasoning: techniques for enterprise systems*. San Francisco, CA: Morgan Kaufmann.

Watson, I (1999) Case-based reasoning is a methodology not a technology. *Knowledge-Based Systems*, **12**, 303-8.

Wood, G D and Ellis, R C T (2003) Risk management practices of leading UK cost consultants. *Engineering, Construction and Architectural Management*, **10**(4), 254-62.