Proceedings of the International Conference on Management of Information Technology for Construction, Singapore, 1993

APPLICATIONS OF MACHINE LEARNING TO CONSTRUCTION SAFETY

TOMASZ ARCISZEWSKI
Associate Professor
Civil and Environmental Engineering Department
Wayne State University
Detroit, Michigan 48202, USA
Affiliated Faculty
Machine Learning and Inference Laboratory
Center for Artificial Intelligence
George Mason University
Fairfax, Virginia 22030, USA.

MUMTAZ USMEN
Professor and chairman
Civil Engineering Department
Wayne State University
Detroit, Michigan 48202, USA.

ABSTRACT

This paper discusses potential applications of machine learning in construction safety. Both learning about accidents and their prevention are described, including examples which demonstrate practical applications of machine learning. Examples were developed using actual accident records and two learning systems: ROUGH, based on the theory of rough sets, and INLEN, based on the STAR methodology. The paper also discusses future research needs and directions.

Key Words
Construction safety; accidents; machine learning; knowledge acquisition; learning about accidents; prevention of accidents.

Introduction

During the recent ASCE conference "Construction in the 21st Century," (Luh-Maan Chang, 1991) the 90's have been declared the decade of construction safety. This reflects a growing realization of the importance of construction safety. Construction has been and continues to be a dangerous occupation, resulting in many accidents, injuries, and fatalities (Hinze and Appelgate, 1991). For example, construction in the USA leads all other industries in OSHA (Occupational Safety and Health Administration) accident incidence rates. In 1990, for instance, construction had 6.7 lost workday cases, 7.5 non-fatal, non-lost workday cases, and 147.9 lost workdays per 100 full-time workers (The term "lost workday" refers to the number of days the worker was not able to report to work because of injury or illness. "Lost workday case" means the number of incidents which resulted in one or more workdays lost.) It has also been reported (ENR, 1992) that construction employs 10 percent of the European Community work force but accounts for 15 percent of the accidents and 30 percent of the fatalities. Therefore, construction safety is a universal problem of a significant importance.

There are several reasons for the construction industry's persistently poor safety record. First, it has a transient work force which is difficult to train. Second, the work involves numerous tasks, procedures, and materials whose handling cannot be standardized. Further, each project is unique in terms of site conditions, plans and specifications, and, to a great extent, contractual arrangements. Third, by and large, the industry lacks the appropriate tools to conduct a sophisticated accident analysis for the development of preventive actions.

Construction safety is important for several reasons. All accidents cause pain and human suffering, and this humanitarian aspect is naturally the most important. However, accidents also have legal and regulatory consequences (i.e., OSHA penalties, liability, and criminal sanctions) and they obviously affect the institutional image for construction companies. Safety also impacts a company's productivity and competitiveness, because companies with better safety records usually provide more efficient work organization and pay smaller workers' compensation premiums. Clearly, accidents and injuries cost money and should be considered in the context of productivity, which is directly affected by their occurrence. Therefore, the construction industry has a lot to gain by improving safety.

It has been shown that a company can control safety by a well-developed safety program (Levitt and Samelson, 1987, Smith and Roth, 1991). A major element of a successful safety program is keeping good records of accidents, so that this information can be utilized for identifying causal factors and developing intervention strategies. At the present time, however, most of the record keeping in the industry is primarily for regulatory compliance purposes, and in larger, more sophisticated firms, accident analyses are performed as a part of loss control/risk management efforts. The analysis methodologies vary; however, they are mostly based on the use of statistical models. Many software packages based on these models are available today. However, most of them are not suitable for use by the smaller companies. On the other hand, as our studies indicate, machine learning about construction accidents could be performed by construction industry personnel at all levels in an easy-to-understand fashion.

We believe that the development of a methodology for machine learning use in construction safety and the preparation of appropriate software, which may result from our work, will fill a large need in the industry. It should be knowledge engineering, dealing with the methodological aspects of using learning systems in engineering. It includes methodologies of evaluation, comparison and selection of learning systems (Arciszewski et al., in print), methodologies of multistage knowledge acquisition (Arciszewski and Mustafa, 1989, Mustafa and Arciszewski, 1989), and methodologies of verification of automatically acquired knowledge. In the area of structural engineering, machine learning is used to learn decision rules in the conceptual design of wind bracings in the steel skeleton structures of tall buildings (Mustafa and Arciszewski, 1992). In transportation engineering, a feasibility study was conducted regarding decision rules for the control of traffic in an urban rail corridor (Khasnabis et al., 1992). In construction engineering, work is concentrated on learning about construction accidents and the development of decision support tools utilizing machine learning for the prevention of accidents (Arciszewski et al., 1991, Usmen and Arciszewski, 1990).

Research on the applications of machine learning to construction safety was initiated by the authors in 1989. Its ultimate objective is to improve construction safety through the prevention of accidents, using enhanced understanding of causal factors affecting accidents and the application of decision support tools for predicting the nature of accidents which might occur under given circumstances. The research has resulted in the development of a set of construction accident descriptors and their nominal values, and in a methodology for acquiring accident data and preparing examples. Also, two feasibility studies of machine learning in acquiring knowledge about construction accidents were conducted. In both case studies the same collection of examples, based on actual accident records provided by Boh Corporation of Louisiana, was used. The first study was performed using the learning system ROUGH (Arciszewski et al., 1991); in the second, INLEN was used (Kaufman et al., 1990). A research plan was also formulated, which includes feasibility studies, learning about construction accidents, and the development of decision support tools for accident prevention. This will utilize the knowledge acquired and will be based on machine learning. In cooperation with the Center for Artificial Intelligence Research of George Mason University, an experimental decision tool for analyzing construction accidents and predicting their nature under given circumstances has been developed. It is based in the AQ15 learning algorithm (Kaufman et al., 1989) and is currently being used for research and demonstration purposes.

Machine Learning

Machine learning is a science dealing with studies and development of computational models of learning and discovery processes and with building learning programs for specific applications. Learning programs, also called "learning systems," are computer programs which transform input in the form of data (usually examples) into knowledge (usually in the form of decision rules). A decision rule is a logical relationship between a group of particularly beneficial for the smaller, less sophisticated firms which are presently not able to perform accident analysis using the statistical approach.

We also predict that in construction safety a paradigm shift will occur: from statistical data analysis to acquiring knowledge about accidents using learning systems. This will be caused by the fundamental advantage of knowledge acquisition based on machine learning with respect to statistical data analysis. In statistical analysis, hypothetical relationships among attributes are produced for the entire population, based on a sample of this population. Therefore this sample must satisfy many specific requirements regarding its size, to obtain reliable results. In machine learning, knowledge is extracted from a given collection of examples, and it is valid in the context of these examples independently of the number of examples used for learning. Obviously, this knowledge can also be considered as hypothetical for the entire population, but it is not so important from the pragmatic point of view.

We foresee that the potential exists for the use of machine learning-based decision support tools in training construction personnel at all levels. These tools can be customized to emphasize economic factors associated with the accidents. They can also be tailored for site supervisors, focusing on specific technical factors. Perhaps these tools could also be adapted for the worker who is interested in correct and safe procedures for a particular type of work. These are factors which led us to begin work on the applications of machine learning to construction safety.

Our research on machine learning in civil engineering was initiated in the Intelligent Computers Laboratory of the Civil and Environmental Engineering Department of Wayne State University in 1985, in cooperation with the Computer Science Department of the University of Regina, Canada. Its initial focus was on knowledge acquisition in structural design (Arciszewski et al., 1987), including methodological aspects, and on the conceptual foundation of using learning expert systems in engineering design (Arciszewski et al., 1987, Arciszewski and Ziarko 1987). At this time, the research concentrated entirely on the applications of learning systems, utilizing learning algorithms based on the theory of rough sets (Pawlak, 1982, Pawlak et al., 1988). Later, this research was expanded to consider other classes of experimental and commercial learning systems such as BEAGLE (Warm Boot, 1988), SuperExpert (Intelligent Terminals, 1987), and a class of systems based on the AQ learning algorithm (Kaufman et al., 1989). At present, our research is carried on in close cooperation with the Computer Science Department of the University of Regina and the Machine Learning and Inference Laboratory of the Center for Artificial Intelligence Research at George Mason University, Virginia, where the first author is an affiliated faculty. Present research concentrates on the development of learning engineering and on the applications of various learning systems in selected areas of civil engineering, including construction, structural, and transportation engineering. Learning engineering is a new domain of accident descriptors (called "independent attributes") and a single accident descriptor (called "a dependent attribute").

After more than fifty years of research, machine learning has reached such a stage of maturity that its engineering applications are not only feasible but also should bring useful results. This is particularly true in the area of construction safety, where no formal mathematical models of accidents are available and the current statistical models are inadequate for several reasons: 1) They require an excessively large number of examples to produce useful results. 2) They are incapable of detecting conceptual patterns or qualitative relationships affecting decision making. 3) Interpretation of the results is difficult. 4) Use of statistical models for making predictions about future accidents is difficult and time-consuming. In addition, machine learning methods can produce useful results from even a small number of examples, and these results can be easily interpreted and linked to human expert knowledge.

Our research is based on a simple paradigm. Records of construction accidents have been accumulated over many years. These records can be used to prepare examples for machine learning, and from these examples knowledge can be extracted. This knowledge will improve our understanding of construction accidents and can be used in various knowledge-based systems, also called "decision support systems," which can be developed and distributed among safety professionals. These systems can be used on an everyday basis to predict accidents and their nature and thus to help prevent them.

Potential Applications of Machine Learning

Two major applications of machine learning in construction safety can be distinguished. First, a learning system can be used as a knowledge acquisition tool to learn about accidents and to acquire knowledge about them. The other application, which may directly affect productivity, is to use a learning system as a decision support tool to prevent accidents through predicting their nature.

Learning about accidents is to be conducted for two reasons: to improve our understanding of accident causal factors and their relationships, and to acquire knowledge for knowledge-based systems. These systems could be used for several purposes, including training, interpretation of accidents, and as decision support tools to prevent accidents. Making predictions about the nature of future accidents is also important, because knowledge of their nature under given circumstances will allow a safety officer to take appropriate preventive actions.

The construction industry is facing a period of intensive change, reflecting progress in various areas of science and technology. It is envisioned that recent developments in computer science, especially in the area of artificial intelligence, will make a major impact on the operations of the construction industry. Inductive and knowledge-based systems may also play an important role in construction safety. In particular, knowledge-based systems containing knowledge about past accidents may become viable safety tools in the near future. It is likely that they will be used by on-site safety personnel on a daily basis for safety management. It is envisioned that the safety officer will consult his/her knowledge-based system to evaluate a safety situation. This consultation will be conducted in the form of a dialogue between the safety officer and a knowledge-based system instaled in a portable computer as demonstrated in the section "**Making Predictions about Accidents**" for the determination of the expected body part injured under given circumstances. Also, inductive tools could be useful for managerial purposes, when economic factors are included in the description of construction accidents and complex planning decisions are considered.

Learning About Construction Accidents

We define learning about construction accidents as four related processes. The first is the elimination of redundant accident descriptors, the second is the introduction of modified attributes which are called "constructed attributes." These two processes can be considered jointly as the modification of the representation space. The third process is the determination of the relative importance of individual attributes, and the fourth is the identification of the logical relationships among various groups of descriptors (attributes), or the extraction of decision rules from examples.

Redundant Attributes

The concepts of redundant attributes and the importance of attributes will be explained in the context of the theory of rough sets (Pawlak, 1982), which has been used as a mathematical foundation for several learning algorithms (Ziarko, 1989a, 1989b). These algorithms have been implemented in various experimental and commercial learning systems and have also been used by us for learning about construction accidents.

An accident can be described by a group of attributes, which can be nominal or numerical. These attributes are divided into independent attributes and a single nominal dependent attribute. The dependent attribute, or decision attribute, is used to classify a given accident into one of several decision categories. The dependency relationship between dependent and independent attributes can be considered in the context of a given collection of examples and measured by the degree of dependency.

The degree of dependency is defined as a percentage of examples which can be classified without any ambiguity into one of the decision categories using the assumed set of independent attributes. When a given independent attribute is eliminated and the degree of dependency remains unaltered, obviously this attribute is redundant and can be eliminated without affecting the dependency. This process of elimination of redundant attributes leads to the determination of a reduct, defined as a minimal set of independent attributes which has the following properties: 1) It preserves the degree of dependency, and 2) No attribute can be eliminated from the reduct without decreasing the degree of dependency. Redundant attributes can be eliminated in various sequences, and therefore for given examples and set of attributes a number of reducts can be produced.

For example, in one of our studies (Arciszewski et al., 1991) a collection of 225 construction accident records was considered and used to prepare examples. A set of thirteen attributes was assumed, and the attribute **Body Part Injured** was selected as the decision attribute. The independent attributes describing the accident victim, the accident, and its results which were used in this study were as follows:

1. Age, 2. Race, 3. Marital Status, 4. Children, 5. Occupation, 6. Job Experience, 7. Time, 8. Season, 9. Accident Type, 10. Work Period, 11. Injury Description, 12. Return to Work.

These attributes and their values were assumed in accordance to the accident records supplied, without any additional studies regarding their nature. They were considered as nominal attributes, and all numerical attributes were converted into nominal attributes. For example, the attribute **Job Experience** (number of years on the job) was converted into a nominal attributes with four values: short, medium, long, very long.

Relative Importance of Attributes

In the case of the attributes considered in our research and the collection of examples used, the degree of dependency was 99.1 percent, i.e., 99.1 percent of examples could be correctly classified into one of the decision categories using the assumed set of independent attributes. The analysis of redundant attributes resulted in the determination of fourteen reducts.

The relative importance of individual attributes from a given reduct can be determined by considering a significance factor. The significance factor indicates the percentage decrease of dependency caused by the removal of a given attribute from the reduct. For example, for the two reducts identified in our research, the values of the significance factor were determined as follows:

Reduct No. 1:

Reduct Attributes	*Significance Factor*
Marital Status	2.7%
Occupation	6.3%
Job Experience	7.3%
Season	1.8%
Accident Type	1.8%
Work Period	3.6%
Injury Description	0.9%

Reduct No. 2:

Reduct Attributes	*Significance Factor*
Race	6.3%
Marital Status	4.5%
Children	2.7%
Job Experience	7.2%
Season	9.4%
Work Period	8.5%
Injury Description	17.5%
Return to Work	1.3%

Individual values of significance factors should be considered in qualitative terms, i.e., absolute values of these factors are much less important than the ratios of these values for individual pairs of factors and attributes associated with them. These ratios can be used to eliminate the least significant but costly attributes and/or to develop a qualitative understanding of the causal accident factors and their expected impact on the accident.

Constructed Attributes

The concept of constructed attributes will be explained in the context of the theory of constructive induction, recently developed at George Mason University (Wnek and Michalski, 1991). Constructive induction is a type of induction in which the formation of a new representation space occurs during inductive learning. This is done by the elimination of some attributes, as in the theory of rough sets-based learning, and by the introduction of constructed attributes. There are two basic types of constructed induction: data-driven and hypothesis-driven. In the first case, various combinations of attributes are considered, using a variety of operations which include addition, subtraction, multiplication, and division of initial attributes. The best combinations of attributes are used in further learning; their selection is based on the performance of the learning system on a given collection of examples. In the case of hypothesis-driven induction, an initial set of classification rules is produced using a learning system based on a selective algorithm (no change in the representation space), and next these rules are used to produce constructed attributes (Wnek and Michalski, 1992).

Logical Relationships among Attributes

The identification of logical relationships between the dependent attribute and various groups of independent attributes is particularly important since these decision rules can be used for numerous purposes, including understanding of accidents and the development of knowledge-based systems. Two examples of decision rules dealing with foot injuries which were produced by two commercial learning systems, ROUGH (Voytec Systems, 1990) and INLEN (Kaufman et al., 1989) are given below. These decision rules were produced using the same collection of 225 examples mentioned in the preceding section. ROUGH produced the following rule:

If **AGE**	is medium, and
RACE	is white, and
SEASON	is first quarter of the year, and
CAUSE	is unsafe act, and
WORK PERIOD	is midday (between the second and the fifth hour of work)

then **BODY PART INJURED** is foot.

A similar decision rule was produced by INLEN:

If **OCCUPATION**	is pile driver, and
INJURY TYPE	is contusion, and
JOB DESCRIPTION	is material handling

then **BODY PART INJURED** is foot.

Both decision rules are relatively simple and can be easily understood by safety personnel who have not been trained in the area of knowledge-based systems.

Machine Predictions About Accidents

Making predictions about accidents is a process in which a learning system is used to predict the values of the decision attribute based on the values of the independent attributes. In this case, the system operates in two stages: learning and consulting. In the learning stage, the system acquires knowledge from examples in the form of decision rules. In the consulting stage, these decision rules are used to make predictions. To illustrate this type, a session with INLEN is briefly described. The objective is to determine the expected value of the decision attribute **BODY PART INJURED** in a case characterized by the combination of independent attributes values, as shown in the answers to individual questions asked by the learning system. The description includes questions, answers, and subsequent predictions produced by the learning system together with the certainty of these predictions. However, in our study we used attributes which had not been specifically developed for the purpose of knowledge acquisition, and therefore might not be adequate for this purpose. For example, one of the nominal values of the attribute **TYPE OF ACCIDENT** is "Hit by foreign matter," and it does not consider the size of this foreign matter. It can be very small, as in the case of eye injuries, or large, as in the case when a worker is hit by a large piece of material and usually sustains spine injuries. This classical artificial intelligence problem of an inadequate representation space in knowledge acquisition is clearly demonstrated in the dialogue given below.

Question No. 1.	What type of accident will it be?
Answer:	Hit by foreign matter.
Preliminary prediction:	No prediction.
Question No. 2.	What type of injury will occur?
Answer:	Foreign body penetration.
Preliminary prediction:	Eye injury (35%) or spine injury (35%).
Question No. 3.	What is the victim's occupation?
Answer:	Carpenter.
Preliminary prediction:	Eye injury (36%) or spine injury (35%).
Question No. 4.	What is the race of the victim?
Answer:	White.
Preliminary prediction:	Spine injury (42%) or eye injury (36%)
Question No. 5.	How much experience does the potential victim have?
Answer:	Over 6 years.
Preliminary prediction:	Spine injury (67%) or eye injury (36%)

Question No. 6.	In what age group is the victim?
Answer:	Over 50 years.
Preliminary prediction:	Spine injury (58%) or eye injury (6%)

Question No. 7.	During what work period will accident occur?
Answer:	First two hours.
Preliminary prediction:	Spine injury (38%) or eye injury (6%)

Question No. 8.	What season will the accident occur in?
Answer:	October to December.
Preliminary prediction:	Eye injury (50%) or spine injury (24%)

Question No. 9.	In what period of the day will the accident occur?
Answer:	Afternoon.
Final prediction:	Eye injury (58%).

In the dialogue reported here, the insufficiently specific attribute values confused the learning system and forced it to conduct a much longer session than would be necessary in the case of properly prepared accident descriptors and their values. To improve the situation, an extensive process of building knowledge representation space would be required, which would include the use of various personal construct-based computer tools for attribute identification. Our session demonstrates, however, the natural character of individual questions, which are easily understandable by non-experts in construction safety and knowledge-based systems.

Making predictions of known examples which have not been used for learning can be also used to evaluate the performance of learning systems. A formal method of such evaluation is given in (Arciszewski et al., 1992). In this method, a formal evaluation procedure, an evaluation model based on the multi-attribute utility theory and a system of fifty evaluation criteria are proposed. As evaluation criteria, various empirical error rates are used, in accordance with Weiss and Kulikowski (1991).

Conclusions

The mathematical theory of machine learning has been under development for the last fifty or so years, and it is reaching the level of maturity where its practical engineering applications are feasible and should bring immediate benefits. This is particularly true for construction safety, where no formal models of accidents are available and any improvement in the understanding of accidents and the development of computer tools for their prevention should be considered as progress with a direct humanitarian and economic impact. Also, recent developments in computer hardware have resulted in inexpensive computers with sufficient computing capabilities to handle advanced machine learning systems and their practical application to construction safety problems.

Research on learning engineering has produced some basic understanding of the methodology of machine learning in engineering, although much more work is still needed. Also, a number of pilot feasibility studies of machine learning in various civil engineering domains has some initial experience.

Machine learning should be given serious consideration as a new technology which may be useful in construction safety. Future research dealing with machine learning in this area should then address two related objectives. The first one is to reformulate construction safety problems and their formal representation into a form (representation space) suitable to learning systems. The second one is to develop a methodological foundation for the use of learning systems in construction safety.

Future research on the methodological foundation of machine learning in construction safety should concentrate on knowledge representation, strategies of automated knowledge acquisition, and methods of formal knowledge verification. Also, after more pilot applications of machine learning in construction safety are completed, medium and full industrial scale applications, leading to the development of knowledge-based decision support tools for accident prevention, should be considered. This research should be accompanied by simultaneous training of safety personnel in machine learning and its safety-related applications.

The research outlined above will be difficult and challenging. It may also be time-consuming, and because of its interdisciplinary nature it may require significant resources. However, considering its expected benefits, saved human lives and construction costs, the challenge of machine learning in construction safety should be met through continued research.

Acknowledgements

The authors are pleased to acknowledge the cooperation of Boh Construction Corporation of New Orleans, Louisiana, and of Reduct Systems, Inc. of Regina, Canada, and of the Center for Artificial Intelligence at George Mason University, Virginia. Boh Construction provided the actual accident data for our experiments. Reduct Systems and the Center for Artificial Intelligence at George Mason University donated the inductive systems ROUGH and INLEN used in these experiments.

The research reported in the paper was conducted in the Intelligent Computers Laboratory of the Civil and Environmental Engineering

Department, Wayne State University, and in the Machine Learning and Inference Laboratory of the Center for Artificial Intelligence, George Mason University. The Center's research is supported by the Defence Advanced Research Projects Agency under Grant No. N00014-91-J-1854, administered by the Office of Naval Research, and Grant No. F90620-92-J-0549, administered by the Air Force Office of Scientific Research, in part by the Office of Naval Research under Grant No. N00014-J-1351, and in part by the National Science Foundation under Grant No. IRI-9020266.

References

Arciszewski, T., Mustafa, M., and Ziarko, W., (1987). "A Methodology of Design Knowledge Acquisition for Use in Learning Expert Systems," Int. J. Man-Machine Studies, Academic Press Limited, 27, 23-32.

Arciszewski T., Ziarko W., (1987). "Adaptive Expert System for Preliminary Engineering Design," Revue Internationale de CFAO et d'Intographie, Vol. 2, No. 1.

Arciszewski, T., and Mustafa, M. (1989). "Inductive Learning Process: The User's Perspective," in Machine Learning, edited by R. Forsyth, Chapman and Hall, 39-61.

Arciszewski, T., Usmen, M., Gleichman, J., (1991). "Construction Accident Analysis: The Inductive Learning Approach," Proceedings of the ASCE Conference "Preparing for Construction in the 21st Century," Boston.

Arciszewski, T., Dybala, T., Wnek, J., (1992). "Method for Evaluation of Learning Systems," Journal of Knowledge Engineering "Heuristics,." Vol. 5, No. 4.

Chang, Luh-Maan, (Editor), (1991). Proceedings of the ASCE Conference, "Preparing for Construction in the 21st Century," Boston.

Hinze, J., Appelgate, L.L., (1991). "Costs of Construction Injuries," Journal of Construction Engineering and Management, Vol. 117, No. 3.

Intelligent Systems (1987). SuperExpert - User Guide, Glasgow, Scotland.

Kaufman, K.A., Michalski, R.S., Kerchberg, L., (1989). "Mining for Knowledge in Databases: Goals and General Description of the INLEN System." Proceedings of IJCAI-89 Workshop on Knowledge Discovery in Databases, Detroit.

Khasnabis, S., Arciszewski, T., Hoda, S.K., (1992). "Learning Rules for Driving Scenarios for an Urban Rail Corridor with Closely Spaced Stations," Proceedings of the ASCE Conference "Computing in Civil Engineering," Dallas.

Levitt, R.L., Samelson, N.M., (1987). Construction Safety Management, McGraw-Hill Book Company, New York.

Mustafa, M., Arciszewski, T., (1989). "Knowledge Acquisition: Engineering Methodology of Inductive Learning," Proceedings of the Workshop on Knowledge Acquisition, International Joint Conference on Artificial Intelligence, Detroit.

Mustafa, M., Arciszewski, T., (1992). "Inductive Learning of Wind Bracing Design for Tall Buildings," Knowledge Acquisition in Civil Engineering, edited by Arciszewski, T., and Rossman, L., ASCE.

Pawlak, Z., (1982). "Rough Sets", Intl. Journal of Computer and Information Sciences, Academic Press Limited, 5(11) 341-356.

Pawlak, Z., Wong, S.K.M., Ziarko, W., (1988). "Rough Sets: Probabilistic versus Deterministic Approach," Intl. Journal of Man-Machine Studies, Vol. 29.

Smith, G.R., Roth, R.D., (1991). "Safety programs and the Construction manager," Journal of Construction Engineering and Management, Vol. 117, No. 2.

Usmen, M., Arciszewski, T., (1990). "Learning from Past Accidents: The Inductive Learning Approach," EXCEL, Newsletter of Center for Excellence in Construction Safety, West Virginia University.

Weiss, S.M., Kulikowski, C.A., (1991). Computers that Learn, Morgan Kaufman Publishers, San Mateo, California.

Wnek, J., Michalski, R., (1991). "Hypothesis-drivem Constructive Induction in AQ17: A Method and Experiments," Proceedings of the IJCAI-91 Workshop on Evaluating and Changing Representations, Sydney, Australia.

Ziarko, W., (1989a). "A Technique for Discovering and Analysis of Cause Effect Relationships in Empirical Data", Intl. Joint Conf. on Artificial Intelligence, Proceedings of the Workshop on Knowledge Discovery in Databases, Detroit.

Ziarko, W., (1989b). "Data Analysis and Case-Based Expert System Development Tool "ROUGH", Proceedings of Case-Based Reasoning Workshop, Pensacola Beach.

Warm Boot Inc. (1987), PC/BEAGLE, User Guide, Nottingham, U.K.

Voytek Systems, (1989). ROUGH - User Manual, Regina, Canada.