STAR METHODOLOGY-BASED LEARNING ABOUT CONSTRUCTION ACCIDENTS AND THEIR PREVENTION

by

T. Arciszewski
R. S. Michalski
T. Dybala

Journal of Construction Automation, Vol. 4, pp. 75-85, 1995.

AUTOMATION IN CONSTRUCTION

Automation in Construction 4 (1995) 75–85

STAR methodology-based learning about construction accidents and their prevention *

Tomasz Arciszewski [a,*], Ryszard S. Michalski [a], Tomasz Dybala [b]

[a] Machine Learning and Inference Center, George Mason University, Fairfax, VA 22030, USA
[b] Computer Science Department, George Mason University, Fairfax, VA 22030, USA

Abstract

This paper presents the results of a feasibility study concerning the application of STAR methodology-based machine learning to construction accidents and their prevention. A ten-stage knowledge acquisition process is presented and its individual stages described. Knowledge about construction accidents was acquired using a collection of 225 examples, based on actual accidents records. Inductive learning with a system based on the STAR methodology was employed. This system was used in both the generalization and specialization modes of operation. The decision rules obtained are complex, but their interpretation is clear and they seem to be consistent with the present understanding of causal relationships between accident results and various factors affecting them. Also, the rules were verified using average overall and omission empirical error rates, which were calculated as average for three randomly determined sequences of examples. These error rates were calculated for all seven steps in the machine learning process, and were used to construct learning curves for both error rates. The relationships between error rates and the number of examples used for learning are analyzed, and coefficients of linear regression given and discussed. The 225 examples used were found to be grossly insufficient to produce reliable knowledge about accidents and therefore a large study is postulated which would involve the collection of a larger number of construction accident records. In general, our study demonstrated the feasibility of machine learning in acquiring knowledge about construction accidents.

Keywords: Construction accidents and their prevention; Knowledge acquisition; Machine learning; Multi-step machine learning process

1. Introduction

Construction accidents cause many human tragedies, are expensive, and disorganize the construction processes [11]. Therefore, their prevention, and even marginal reductions in their cost, will have a significant human and financial impact. Prevention of construction accidents usually requires predicting future accidents and their nature under given circumstances [12,14]. Making such predictions must be based on knowledge about past accidents, and can be conducted using various decision support tools, including those utilizing machine learning. For these reasons, in

* Discussion is open until October 1995 (please submit your discussion paper to the Editor on Construction Technologies, M.J. Skibniewski).

* Corresponding author.

0926-5805/95/$09.50
SSDI 0926-5805(94)00035-2

1989, research on automated knowledge acquisition in learning about construction accidents and their prevention was initiated in the Intelligent Computers Laboratory of the Civil Engineering Department, Wayne State University, Detroit. The initial research was on the application of the rough sets-based machine learning, and concentrated on learning about construction accidents and required methodology. The results were reported in [4]. The present research is focused on the methodologies of both learning about construction accidents and their prediction. The second component is particularly important. The methodological foundation for learning about construction accidents has already been developed, and can be used for practical purposes. However, learning about accidents may only indirectly improve construction safety, while the development of decision support tools for predicting future accidents and thus helping prevent them should bring direct and immediate safety gains.

The objective of this paper is to investigate the feasibility of STAR methodology-based learning in knowledge acquisition about construction accidents and their prevention. The paper presents the results of a feasibility study conducted in the Machine Learning and Inference Laboratory of the Center for Artificial Intelligence at George Mason University, Fairfax, Virginia. The examples of accidents were prepared in the Intelligent Computers Laboratory, Civil Engineering Department, Wayne State University, Detroit, Michigan. Four major questions were addressed:

1. How difficult is to prepare examples of construction accidents for machine learning-based knowledge acquisition and prediction of future accidents, and how can this be done?
2. What knowledge in terms of decision rules could be expected?
3. How truthful is knowledge acquired from a relatively small number of examples?
4. How many examples are necessary to obtain knowledge and to make predictions about future accidents with high reliability?

This paper provides the description of the entire process of knowledge acquisition, including a multi-step machine learning process. The learning system used is also briefly described. The decision rules obtained are demonstrated, discussed, and verified using the overall and omission empirical error rates. The relationships between the error rates and the number of examples are also examined. Major conclusions are given and directions for further research discussed. We postulate using a much larger number of examples to obtain results of practical significance.

2. Knowledge acquisition

Knowledge acquisition about construction accidents was conducted as a formal process, which was developed for the purposes of the research reported here. This process and its principal component, a multi-step machine learning process, are defined as follows:

The knowledge acquisition process is the entire process of transforming input in the form of construction accidents data into output in the form of decision rules.

The multi-step machine learning process is a part of the knowledge acquisition process in which a learning system is used to transform input in the form of examples of construction accidents into output in the form of decision rules.

A ten-stage knowledge acquisition process with the following stages was assumed [6]:

1. Collection of accidents records.
2. Identification of accident descriptors (attributes).
3. Initial preparation of examples.
4. Initial automated knowledge acquisition.
5. Analysis of the knowledge acquired.
6. Modification of attributes and clustering of their nominal values.
7. Final preparation of examples.
8. Final automated knowledge acquisition.
9. Analysis of decision rules and their interpretation.
10. Formal verification of knowledge acquired.

A total of 225 accident records were provided by Boh Corporation, Louisiana (Stage 1). This number of examples is small for machine learning, particularly considering the complexity of the

Table 1
Attribute *Job Experience*: conversion table

Numerical value	Nominal value
0 < Job Experience < 6 months	short
6 months < Job Experience < 2 years	medium
2 years < Job Experience < 6 years	long
Job Experience > 6 years	very long

problem of construction accidents. The entire representation space contains $4.2024 \cdot 10^9$ possible events (the number of all possible events is calculated by multiplying the numbers of values for each attribute), and the examples provided are only 0.0000054 percent of this space. However, collecting accident records proved to be much more difficult and time-consuming than expected. Therefore, it was decided that our initial collection of 225 accident records would have to be sufficient for the feasibility study planned, which was never intended to produce knowledge for practical applications.

The accident records were used to prepare the collection of examples. First, the records were analyzed and the initial set of construction accident descriptors, called attributes, was identified (Stage 2). The attributes were both nominal (for example: *Marital Status*) and numeric (for example: *Number of Children*). Next, all numeric attributes were converted into nominal form using conversion tables. An example of such a conversion table is shown in Table 1 for the attribute *Job Experience*. All attributes and their nominal values were used to prepare the initial collection of examples (Stage 3). This collection is most likely noisy in terms of incorrectly interpreted and recorded values of individual attributes, but it is a realistic example of construction data available for automated knowledge acquisition.

The collection of examples was used in the initial automated knowledge acquisition process (Stage 4), which was conducted with the learning system ROUGH, developed by Voytech Systems, Inc. [17]. ROUGH is based on the theory of rough sets [13,18]. This, like the other automated knowledge acquisition processes reported in this paper, was conducted as a multistage process, and learning was done in several stages [5,10]. The entire collection of available examples was divided into groups related to individual stages. These groups were gradually added, in the subsequent stages, to the initial group of examples selected for the first stage. At each stage of the process, the learning system was used to produce decision rules for a given group of examples.

ROUGH produced 146 decision rules for the attribute *Body Part Injured*, which was selected as the decision attribute. The decision rules were analyzed in cooperation with an experienced construction safety professional (Stage 5) and were found to be unsatisfactory because of their excessive specificity, which reflected the nature of the examples in terms of attributes and their specific values assumed in accordance with the construction accident recording guidelines used by Boh Corporation. Therefore, we decided that the representation space must be reduced. All attributes and their values were reexamined, with two objectives [2,3]: (1) to eliminate attributes which were irrelevant to the construction accidents represented by the examples, and (2) to cluster the values of individual attributes. As the result of this analysis, the number of attributes was reduced from twenty to thirteen. Such attributes as *Sex*, *Wage per Hour*, *State*, *Accident Place*, *Accident Cause*, *Disease*, and *Fatal Result* (a binary attribute) were eliminated. Finally, the following four groups of attributes were assumed:

I. **Personal Information**
   1. *Age*
   2. *Race*
   3. *Marital Status*
   4. *Children*

II. **Job Information**
   5. *Occupation*
   6. *Job Experience*

III. **Accident Information**
   7. *Hour of Day*
   8. *Season*
   9. *Work Period*
   10. *Accident Description*

IV. **Injury Information**
   11. *Injury Type*
   12. *Body Part Injured*
   13. *Return to Work*

An example of clustering of attribute values for the attribute *Marital Status* is shown in Table

Table 2
Attribute *Marital Status*: clustering of values

Initial values	Final values
married	married
single	unmarried
divorced	unmarried
separated	unmarried
widower	unmarried
unknown	unknown

2. The left column shows the initial values, the right one the final values.

Next, final attributes and their new values were used (Stage 7) to prepare the final collection of examples. For this collection, two groups of examples were subsequently prepared for two different attributes considered as decision attributes, namely (1) *Body Part Injured*, and (2) *Job Experience*.

The first collection of examples represents input for machine learning to produce output in the form of decision rules. These rules relate the decision attribute *Body Part Injured* to the remaining attributes, which in this case are considered as independent attributes.

The second collection of examples is input for machine learning algorithm to produce decision rules which relate *Job Experience* to the remaining attributes. To illustrate the nature of the examples, one of their instances for the decision attribute *Body Part Injured* is shown below:

Example 67:
Job Experience = over six months
Age = 30 to 50 years
Race = Nonwhite
Marital Status = married
Children = yes
Occupation = laborer
Hour of Day = AM
Season = January to March
Accident Description = striking object
Work Period = 2 to 6 hour
Injury Type = dislocation
Return to Work = yes
the decision attribute Body Part Injured = leg injury

The decision attributes chosen were selected for the two reasons: (1) to investigate two extreme and most interesting cases of automated knowledge acquisition, and (2) because of their domain significance. In this way, we intended to develop an understanding of machine learning over the entire spectrum of construction accidents problem in terms of the nature of attributes and the number of their nominal values.

In the first case, the decision attribute *Body Part Injured* can attain one of seventeen nominal values, and the numbers of examples for the individual categories are relatively small. This can be called a "weak attribute". In addition, in construction safety it is an important attribute, because the cost of an accident strongly depends on the part of the body injured. Therefore, the ability to predict what part of the body will most likely be injured under given circumstances provides an opportunity to take preventive measures and eventually to avoid a given accident.

The second decision attribute investigated was *Job Experience*. It can attain only five nominal values, and the numbers of examples for the individual categories are relatively large. This can be called a "strong attribute". This attribute is also quite important, because the ability to predict if a certain accident will occur depending on the job experience of a given worker may lead to accident avoidance through the replacement of the worker by a more experience person, if necessary.

The results of learning were investigated in Stage 9, which is reported in Section 4. Finally, the knowledge produced was verified in Stage 10, described in Section 5.

3. The STAR methodology

The results presented in this paper were obtained using INLEN (Inference and Learning), a computer software package for automated rule learning and building decision support tools. INLEN was developed in the Center for Artificial Intelligence at George Mason University [9]. It is based on the learning algorithm AQ15, originally proposed by Michalski [8]. The AQ15 learns classification rules from training instances consisting of sample patterns and their correct classification. The algorithm seeks to find the most general rule in the rule space that discriminates training instances in class c_i from all training instances in all other classes c_j $(i \neq j)$. The learned rules are called discrimination rules. The representation language used in AQ15 is VL_1, an extension of the propositional calculus. VL_1 is a fairly rich language that includes conjunction, disjunction, and a set-membership operators. Consequently, the rule space of all possible VL_1 discrimination rules is quite large.

The algorithm uses an inductive learning method utilizing STAR methodology [7]. *A STAR of the event e against the event set E* is defined as a set of all maximally general conjunctive expressions that cover event e and that do not cover any of the examples in set E, where e (an event) is a positive example of a concept to be learned and E is a set of some counter examples of this concept. In practical problems a star of an event may contain a very large number of descriptions. Consequently, such a theoretical star is replaced by *bounded star* that contains no more than a fixed number of descriptions. These descriptions are selected as the most preferable descriptions, according to the preference criterion defined in the problem background knowledge.

A general algorithm utilizing STAR methodology can be described as follows [7]:

Step 1. Randomly select a positive example.
Step 2. Generate a bounded star of that example against the set of negative examples. In the process of star generation apply generalization rules, task specific rules, heuristic for generating new descriptors supplied by problem background knowledge, and definitions of previously learned concepts.
Step 3. In the obtained star, find a description with the highest preference according to the assumed preference criterion.
Step 4. If found description covers set of positive examples completely, then go to Step 6.
Step 5. Otherwise, reduce the set positive examples to contain only events not covered by learned description and repeat the whole process from Step 1.
Step 6. The disjunction of all generated descriptions is complete and consistent concept description. As a final step apply various reformulation rules defined in background knowledge in order to obtain simpler expression.

The central step in the above methodology is the generation of a bounded star. This can be done using a variety of methods. Thus, the above STAR methodology can be viewed as a general schema for implementing various learning methods and strategies. Details about generation of bounded star are described in [7].

The AQ15 algorithm which directly employes STAR methodology, can be tuned using many input parameters which control the execution of the program. The authors investigated two possible modes of operation, generalization and specialization. The generalization mode induces rules as general as possible, i.e., they involve the minimum number of extended selectors, each with the maximum number of values. The specialization mode generates rules as simple as possible, i.e., with the maximum number of extended selectors and the minimum number of values. Redundant values are removed from extended selectors in the rule [8].

4. Knowledge acquired

Automated knowledge acquisition was conducted to produce decision rules relating individual decision attributes to the remaining attributes. For each decision attribute, the learning system was used in both the generalization and specialization mode. In this way, four collections of decision rules were produced. All decision rules were generated in the disjunctive normal form [8], also called the standard normal form [16] with internal disjuncts of attribute values, i.e., each collection of decision rules can be considered as a set of independent rules, and each rule represents a set of conditions which must be satisfied. In addition, conditions for individual may attain only one value from a given set of attribute values, and the power of this set may be equal to or greater than one. The use of decision rules in the standard normal form with internal disjuncts led to the generation of a much smaller number of decision rules than would otherwise have been produced.

To illustrate the nature of the decision rules produced, two simple but complete sets of decision rules relating head injuries (*Body Part Injured* = Head) to the remaining attributes are given for both the generalization and specialization modes of learning.

For example, in the generalization mode decision rules like the following were obtained:

1. ***Head Injury should be expected in an accident when:***

1. *Occupation* of victim is **laborer or pile driver.**
2. *Job Experience* is **short** or **very long.**
3. *Hour of Day* is **morning.**
4. *Season* is **Summer** or **Fall.**
5. *Accident Description* is **injuries by lifting, pulling, pushing** or **injuries while handling material** or **striking against various objects.**
6. *Injury Description* is **contusion, laceration, puncture, foreign body struck** or **fracture.**

or

2. ***Head Injury should be expected in an accident when:***

1. *Race* of victim is **nonwhite.**
2. *Occupation* is **carpenter** or **pile driver** or **support staff.**
3. *Job Experience* is **medium** or **long**
4. *Season* is **Spring** or **Fall.**
5. *Accident Type* is **injuries while handling material** orslips **and falls.**
6. *Injury Description* is **contusion, laceration, puncture** or **fracture.**

When the specialization mode was used, two decision rules were also produced:

1. ***Head Injury should be expected in an accident when:***

1. *Occupation* of the victim is **pile driver** or **support staff.**
2. *Accident Description* is **material handling** or **slip / fall.**
3. *Season* is **spring** or **fall.**
4. *Injury Type* is **contusion** or **fracture.**
5. *Victim* has **children.**
6. *Work Period* is **not the first two hours of work.**
7. *Marital Status* is **married.**
8. *Job Experience* is **between medium or long.**
9. *Age* is **under 50 years.**

or

2. ***Head Injury should be expected in an accident when:***

1. *Injury Type* is **fracture.**
2. *Season* is **Fall.**
3. *Accident Description* is **material handling.**
4. *Job Experience* is **medium.**
5. *Victim* has **Children.**
6. *Work Period* is **not the first two hours.**
7. *Marital Status* is **married.**
8. *Hour of Day* is **morning.**
9. *Age* is **less than 50 years old.**
10. *Return to Work* is **with no lost time.**
11. *Occupation* is **carpenter** or **laborer.**

The decision rules presented here are the sim plest among those produced. However, even these rules are too complex to be produced manually. They clearly show why the present understanding of construction accidents is still so limited, and demonstrate the superiority of the automated knowledge acquisition over to any other form of learning. It should be noted, however, that these decision rules are valid only in the context of the collection of examples used to produce them, and that they are only plausible hypotheses for the entire representation space [15].

5. Knowledge verification

One of our objective was to investigate the feasibility of STAR methodology-based learning in construction accident prevention. It was assumed that the ability to prevent accidents is dependent on the ability to predict them. Therefore, the feasibility of STAR methodology-based learning in accident prevention has been studied based on the analysis of the accuracy of predictions about unseen accidents, based on the knowledge acquired from known accidents. This accuracy can be formally measured by various empirical error rates, which are determined through tests. In each test, a learning system uses a given body of examples to make predictions about other known examples which have not been included in its training input set. Each test can then be compared to a real-life situation, when a construction safety professional working on the prevention of accidents uses a decision support system to predict future possible accidents and their nature. Therefore, empirical error rates are highly relevant to both machine learning research, which is concerned with the performance of learning systems, and to construction safety, which is concerned with the prevention of accidents and improvement of safety.

Knowledge verification was conducted as part of the process of automated knowledge acquisition. This was done using the method proposed in [19]. In our verification, empirical error rates were determined at the individual stages of the automated knowledge acquisition process, and learning curves were constructed for these rates. A learning curve is understood as a graphical relationship between a given error rate and the number of examples used to generate the decision rules utilized to conduct the tests and to produce this error rate.

Two error rates were used: (1) the overall empirical error rate, and (2) the omission error rate. The overall empirical error rate was used because it provides the most general evaluation of performance of a learning system and the knowledge acquired. Also, it has a simple interpretation, convincing to construction safety professionals. The omission empirical error rate is also important, because it measures the degree to which the learning system, using the knowledge acquired, fails to recognize cases belonging to individual categories of the decision attribute.

The overall empirical error rate is defined [1] as

$$E_{oe} = \frac{\text{number of errors}}{\text{number of tests}},$$

where

an error = a misclassification of a test example,

number of tests

= number of classification tests (predictions).

The omission empirical error rate is defined [1] as

$$E_{om} = \frac{\sum_{i=1}^{n} E^i_{om}}{n},$$

where

n = the number of classes,

$$E^i_{om} = \frac{\text{number of omission errors for class } i^*}{\text{number of tested positive examples of class } i^*}$$

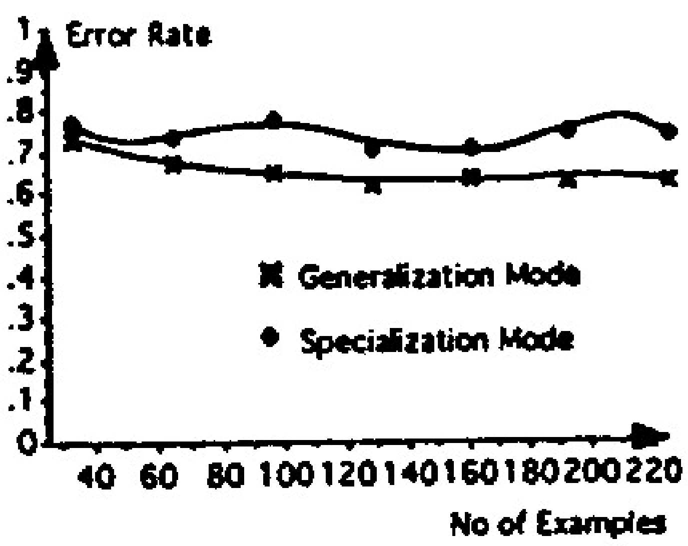

a. Learning Curves

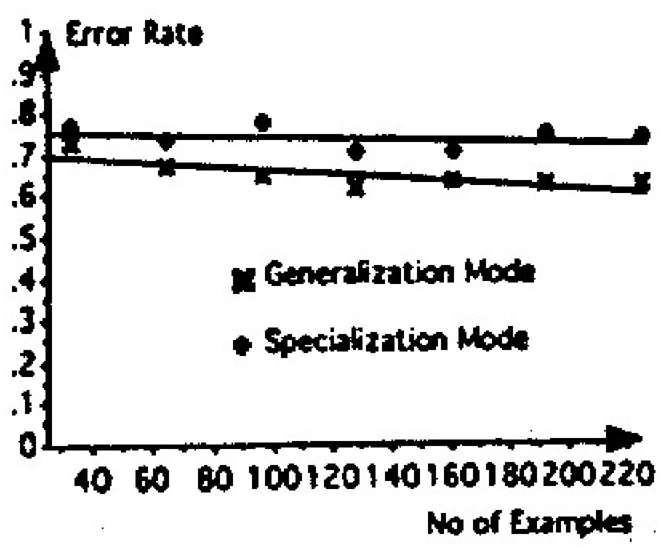

b. Linear Regression Lines

Fig. 1. Overall empirical error rate: decision attribute *Body Part Injured*.

with

number of omission errors for class i''
= number of errors when a positive example is classified as a negative one,

number of tested positive examples for class i''
= number of classification tests for class i'' examples.

For each decision attribute considered and each mode of operation of the learning system, both empirical error rates were calculated for all seven stages of the automated knowledge acquisition process, using the leave-one-out resampling method [16] for three different sequences of examples and the final results were produced as the average empirical error rates for these three sequences.

Figures 1(a) and 2(a) show the learning curves for the overall average empirical error rates for decision rules produced for the decision attributes *Body Part Injured* and *Job Experience*, obtained for both the generalization and specialization modes of operation.

In Figs. 1(b) and 2(b), linear regression lines are also given to show the nature of the changes in the error rates which occur throughout the automatic knowledge acquisition process. In most cases, the error rates decline with the growing number of examples, and the average reduction of empirical error rates is 5.2 percent for the entire process. In only two cases, empirical error rates actually increased with the number of examples, but the increase is relatively small, 4.6 to 4.7 percent. In both cases the learning system was used in the specialization mode of operation. The increase occurred for the omission error rate for the decision attribute *Body Part Injured* and for the overall error rate for the decision attribute *Job Experience*. This can be explained by the small number of examples considered, when compared to the total number of possible events in the representation space, and by the fact that significant changes in the error rates, both increases and declines, usually occur at the beginning of learning.

Table 3
Empirical average error rates: final values

Decision attribute	Overall error rates		Omission error rates	
	Special	General	Special	General
Body Part Injured	73%	63%	84%	76%
Job Experience	78%	59%	80%	64%

Table 3 provides final numerical results for all average empirical error rates. A significant difference (9 to 24 percent) in performance can be observed between the specialization and generalization modes. In the generalization mode, the system produces the most general complexes [8] which cover larger numbers of unseen events in the representation space than the more specific complexes produced in the specialization mode.

Differences in empirical error rates occur also between corresponding rates for the two decision attributes considered. In general, the performance of the learning system is better (on aver-

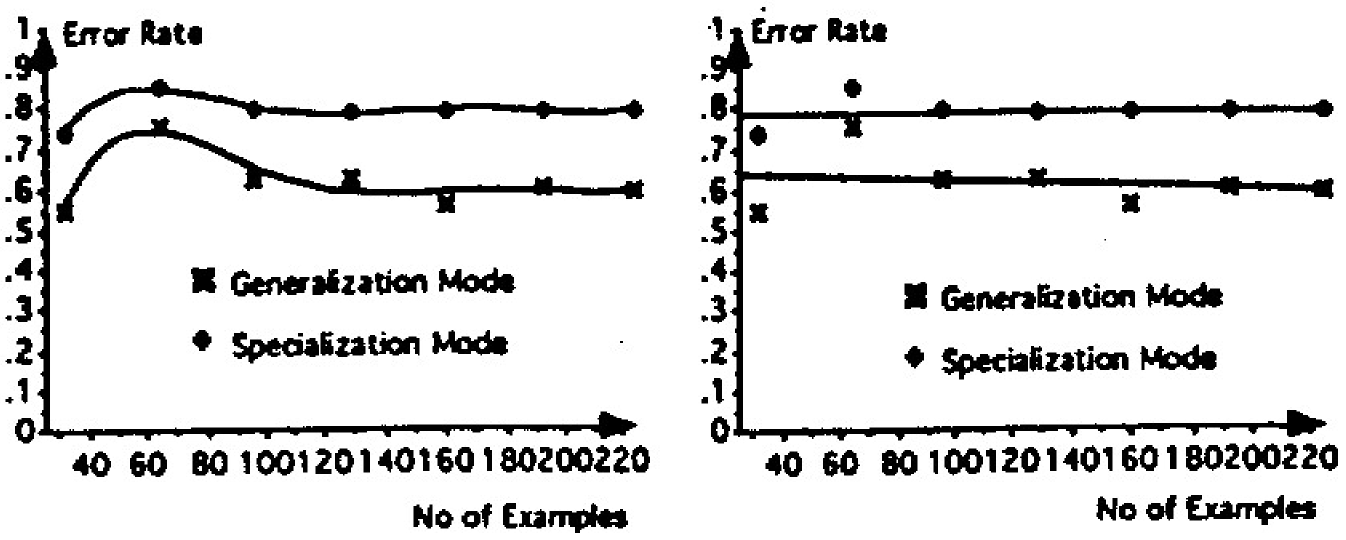

a. Learning Curves b. Linear Regression Lines

Fig. 2. Overall empirical error rate: decision attribute *Job Experience*.

age, 5 percent) on the decision attribute *Job Experience* than on the attribute *Body Part Injured*. This can be explained by the difference in sizes of the representation spaces: the representation space is smaller for the decision attribute *Job Experience*, and therefore the performance of the learning system is better. For the overall average error rate in the specialization mode, a better result is obtained for the decision attribute *Job Experience*, but this exception can be explained as before by the relatively small size of the collection of examples.

6. Number of examples versus error rates

As discussed in Section 2, collecting examples is difficult and expensive. For this reason, a determination of the relationships between the number of examples and empirical error rates is essential, and therefore significant attention was paid to this problem. The relationship sought can be used for two purposes: (1) to monitor the progress of the automated knowledge acquisition process, and (2) to determine how many examples are required in a given case to make predictions about future accidents with an assumed error rate.

All the average empirical error rates, discussed Section 4, were used in a linear regression analysis to produce the functional relationships between the number of examples (independent variable) and the empirical error rates (dependent variables). The individual linear function coefficients, their slopes and constants, are shown in the Table 4. The linear regression lines are shown in Figs. 1(b) and 2(b).

The linear regression relationships obtained are formally valid only within the range 0-225 accident examples, and any extrapolation of these relationships may produce invalid results. However, using these relationships and making an additional assumption about their linearity outside the range 0-225, it could be speculated that approximately 1100 examples would be necessary to produce decision rules enabling the learning system to make predictions about *Body Part Injured* in future accidents with an overall error rate of 20 to 30 percent. Such error rate is typical for experienced human experts, and the availability of a decision support system with similar performance would be a significant development, particularly in that there are only few human experts in the area of construction accidents, and the collection of 1100 examples could easily be accomplished.

7. Conclusions

The results presented in this paper were produced by a short-term feasibility study. These results were obtained using only a single learning system, INLEN, which is based on the STAR learning methodology. Other learning systems might produce different results, but our experience in the evaluation and comparison of various systems indicates that no significant difference is to be expected, particularly when similarly advanced learning systems are used.

Table 4
Empirical error rates: linear regression coefficients

Decision attribute	Mode	Error rate	Slope	Constant
Body Part Injured	Genaralization	Overall	$-4.68\text{E}-4$	0.711
		Omission	$-1.18\text{E}-4$	0.788
	Specialization	Overall	$-1.59\text{E}-4$	0.760
		Omission	$1.68\text{E}-4$	0.810
Job Experience	Generalization	Overall	$-1.22\text{E}-4$	0.649
		Omission	$-2.41\text{E}-4$	0.725
	Specialization	Overall	$1.65\text{E}-7$	0.789
		Omission	$-4.79\text{E}-6$	0.807

The developed ten-stage knowledge acquisition process was adequate for the purpose of learning about construction accidents. However, collection of accident records, identification of attributes, and preparation of examples were found much more difficult and time-consuming than expected. Also, significant effort was required to modify individual attributes and to cluster their values. This work was chiefly necessary because it was the first application of machine learning in the construction accident area, and no methodological experience was available when the research was initiated.

INLAN was used in both the generalization and specialization modes of operation, and significant differences in results were observed. In general, INLAN in the generalization mode produced better results in terms of empirical error rates than in the specialization mode. The average difference is 17 percent with respect to the error rates for the generalization mode. Also, the numbers of decision rules were different, for example, 34 versus 32 for the decision attribute *Body Part Injured*.

The decision rules which were generated are complex, but their interpretation is clear and they seem to be consistent with the present understanding of causal relationships between accident results and various factors affecting construction accidents. However, these decision rules were produced from a very small number of examples considering the size of the representation space, and therefore they should not be used for any practical purposes. They are valid only in the context of the actual examples used for their generation. For all other cases, these decision rules are only weak plausible hypotheses about construction accidents, as was demonstrated in Section 5.

The knowledge produced was verified using the overall and omission empirical error rates and considering seven-step learning processes. For these processes, learning curves were constructed for both error rates and for both modes of operation. Empirical error rates are high, in the range of 60 to 85 percent for the final step of the learning process. This disappointing but expected result indicates the insufficient number of examples used for learning, particularly considering the complex nature of construction accidents reflected in the size of the representation space. The other result supporting this observation is that two error rates out of the eight considered actually increased during the seven-step learning process. This is typical for the beginning of learning.

Our study demonstrates the feasibility of machine learning to acquire knowledge about construction accidents and their prevention. Considering the importance of construction accidents and their high social and economic costs, a full-scale investigation is justified. Such an investigation should be conducted using a much larger number of construction accident examples, and preferably using several systems based on various learning algorithms. Only such an investigation could produce results with a practical impact. All the necessary theoretical and methodological background for such investigation has already been prepared, and realization is a matter of finding sufficient resources.

Acknowledgements

The authors would like to acknowledge the contribution of Dr. Muntaz Usmen, Civil Engineering Department, Wayne State University, who prepared construction accident descriptors, and of Gamal Arafat, Peter Hajdo, and Khalid Nasim, Ph.D. students in the same department who prepared the examples of accidents. This research was conducted in the Center for Artificial Intelligence at George Mason University. The Center's research is supported in part by the National Science Foundation under the grant No. IRI-9020266, in part by the Advanced Research Projects Agency under the grant No. N00014-91-J-1854, administered by the Office of Naval Research, and under the grant No. F49620-92-J-0549, administered by the Air Force Office of Scientific Research, and in part by the Office of Naval Research under the grant No. N00014-91-J-1351.

References

[1] T. Arciszewski, T. Dybala and J. Wnek, Method for evaluation of learning systems, *J. Knowledge Engrg. Heuristics*, *2* (4) (1992).

[2] T. Arciszewski and L. Rossman (Eds.), *Knowledge Acquisition in Civil Engineering*, American Society of Civil Engineering, New York (1992).

[3] T. Arciszewski and W. Ziarko, Machine learning in knowledge acquisition, in: T. Arciszewski and L. Rossman (Eds.), *Knowledge Acquisition in Civil Engineering*, American Society of Civil Engineers, New York (1992).

[4] T. Arciszewski, M. Usmen and J. Gleichman, Construction accident analysis: the inductive learning approach, *Proc. ASCE Construction Congress Construction in the 21st Century*, Boston, April 1991.

[5] T. Arciszewski and M. Mustafa, Inductive learning process: the user's perspective, in: R. Forsyth (Ed.), *Machine Learning*, Chapman and Hall, London (1989) 39–61.

[6] T. Arciszewski, M. Mustafa and W. Ziarko, A methodology of design knowledge acquisition for use in learning expert systems, *Int. J. Man-Machine Studies*, *27* (1987) 23–32.

[7] R.S. Michalski, Theory and methodology of inductive learning, in: R.S. Michalski, J.G. Carbonell and T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach*, Tioga Publishing, Palo Alto, Calif. (1983).

[8] R.S. Michalski, I. Mozetic, J. Hong and N. Lavrac, The AQ15 inductive learning system: an overview and experiments, George Mason University, Reports of Machine Learning and Inference Laboratory, No. MLI-86-6, 1986.

[9] R.S. Michalski, L. Kerschberg, K. Kaufman and J. Ribeiro, Mining for knowledge in databases: the INLEN architecture, initial implementation and first results, *Int. J. Intelligent Systems*, *1* (1992).

[10] M. Mustafa and T. Arciszewski, Knowledge acquisition: engineering methodology of inductive learning, *Proc. Workshop on Knowledge Acquisition, International Joint Conf. on Artificial Intelligence*, Detroit, 1989.

[11] National Academy Press, *Injury in America, A Continuing Public Health Problem*, Washington, D.C. (1985).

[12] NIOSH, National traumatic occupational safety fatalities 1980–1984, Report by NIOSH Division of Safety Research, Morgantown, West Virginia, 1987.

[13] Z. Pawlak, Rough sets, *Int. J. Comput. Inform. Sci.*, *5* (11) (1982) 341–356.

[14] E.S. Pollack and D.G. Keimig (Eds.), *Counting Injuries and Illnesses in the Workplace: Proposal for a Better System*, Report, National Academic Press (1987).

[15] M. Usmen and T. Arciszewski, Learning from past accidents: the inductive learning approach, *EXCEL* (Newsletter of the Center for Excellence in Construction Safety), West Virginia University (1990) 11–12.

[16] S.M. Weiss and C.A. Kulikowski, *Computers that Learn*, Morgan Kaufman, San Mateo, Calif. (1991).

[17] Voytec Systems Inc., ROUGH – User's Manual, Regina, Canada, 1990.

[18] W. Ziarko, Data analysis and case-based expert system development tool ROUGH, *Proc. Workshop on Case-Based Reasoning*, Pensacola Beach, Fla., Morgan Kaufman, San Mateo, Calif. (1989).

[19] W. Ziarko and T. Arciszewski, Verification of morphological table based on probabilistic rough sets approach, *J. Particulate Sci. Technol.*, 6 (1988) 193–205.