94-19

MACHINE LEARNING IN TRANSPORTATION ENGINEERING: A FEASIBILITY STUDY

T. ARCISZEWSKI
Systems Engineering Department
School of Information Technology and Engineering
George Mason University, Fairfax, Virginia, USA

S. KHASNABIS
Department of Civil Engineering and
Urban Transportation Institute
Wayne State University
Detroit, Michigan, USA

S. KHURSHIDUL HODA
Carl Walker Engineers
Kalamazoo, Michigan, USA

W. ZIARKO
Computer Science Department
University of Regina
Regina, Canada

This paper presents the results of a feasibility study on the application of machine learning to knowledge acquisition in transportation engineering. An eight-stage knowledge acquisition process is proposed and its individual stages justified and described. Machine learning is used to learn about urban rail control. The development of the representation space for this problem is discussed, including the analysis of motion and stopping regime for a train, and of both decision and performance attributes. Travel time, energy consumption, and passenger comfort are used as performance attributes. Six automated knowledge acquisition processes were conducted for various performance (dependent) attributes, taking into consideration two different clusterings of performance attribute values into three and seven subranges. All the examples used for learning were computer generated, using REGIME, which separately produces estimations of individual performance attributes for a given train-driving scenario and an assumed rail corridor. The decision rules produced are discussed, and their verification, based on the overall empirical error rate, is reported. This paper also contains conclusions and suggestions regarding future research on applications of machine learning to knowledge acquisition in transportation engineering.

The study was conducted during the 1990–92 period at Wayne State University and was sponsored by the U.S. Department of Transportation through the IVHS Program at the University of Michigan.

Address correspondence to T. Arciszewski, Associate Professor of Urban System Engineering, Systems Engineering Department, School of Information Technology and Engineering, George Mason University, Fairfax, VA 22030, USA.

Applied Artificial Intelligence, 8:109–124, 1994

0883-9514/94 $10.00 + .00

Transportation engineering is a subarea of civil engineering concerned with various transportation systems and their modeling, design, and maintenance, including physical maintenance and the control of vehicle flow. This domain can be characterized by three major features: a large number of complex problems whose solution requires the use of complicated mathematical models, the probabilistic nature of transportation phenomena, and the availability of records of past events: accidents, traffic flow patterns, traffic control scenarios, etc. Because of complexity of the formal models in transportation engineering, their use is difficult and time consuming. Therefore, the classical deductive engineering approach to problem solving is often of limited value. The availability of examples in such cases makes the inductive approach particularly attractive, especially when it is based on the use of machine learning and on knowledge produced by knowledge-based decision support systems.

The long-term objectives of our research are to investigate and to determine the feasibility of various forms of machine learning in transportation engineering and to explore the advantages and disadvantages of the individual learning paradigms and learning systems on which they are based. In this paper, we present an improved understanding of the methodological aspects of automated knowledge acquisition in transportation, and we report some new domain knowledge produced as the result of our research. The experiments were conducted in the area of traffic control in an urban rail corridor with closely spaced stations. This study is designed to supplement research on intelligent vehicles highway systems (IVHS) in order to demonstrate the feasibility of machine learning in knowledge acquisition for control of intelligent vehicles.

The concept of preprogrammed driving for urban rail corridors has been proposed in the literature. It will permit an automated selection of driving scenarios consistent with the distribution of ridership demand along the corridor. The driving scenarios are likely to change as demand changes with the time of the day. This concept is consistent with IVHS technology that aims at the integration of the vehicle, the facility, and the driver using state-of-the-art communication, computer, and electronic technology (Mobility 2000 Working Group, 1990; Underwood et al., 1989).

PROJECT JUSTIFICATION

Transportation engineering is a rapidly changing area, characterized by the development of new transportation systems and technologies. Particularly important is the current research on intelligent vehicles technology, which includes development of new methodologies of traffic control for both road and rail vehicles. As mentioned earlier, the direct use of formal mathematical models in traffic control is difficult because of their complexity and the time required to produce results. For these reasons, a knowledge-based approach to traffic control is considered a promising alternative. However, this approach requires the availability of formal knowledge in a form suitable to knowledge-based systems, preferably decision rules. Since traditional manual methods of knowledge acquisition are unreliable in the case of complex engineering problems (Modesitt, 1992), the determination of the feasibility of automated knowledge acquisition was crucial to progress in intelligent vehicles research. The actual engineering needs to be satisfied have thus led us to consider machine learning and to explore various learning systems in the automated acquisition of knowledge about urban rail driving scenarios.

Our analyses were conducted in the area of traffic control of an urban rail corridor with closely spaced stations. This specific area was selected for several reasons. First, it is a relatively well-understood domain, and results of automated knowledge acquisition could be verified by human experts. Second, it was feasible to develop formal mathematical models describing traffic in such corridors; this was necessary to prepare examples for machine learning experiments. Third, IVHS technology is expected to be used for rail vehicles, and therefore, we wanted our research to be concentrated in an area of practical impact.

A learning system based on rough sets was chosen for our research for several reasons. First, the theory of rough sets provides mathematical models of imprecise knowledge representation, analysis, and acquisition developed with full formal rigor by logicians and computer scientists (Pawlak, 1992). This makes the theory and its learning, or analytical methods, well understood and traceable in terms of formal, unbiased reasoning methods. The method of rough sets in its application to classification problems has also been proven in a number of areas ranging from medical diagnosis to process control (Pawlak et al., 1992; Slowinski, 1992). These sound theoretical foundations, combined with promising existing applications and the availability of commercial software packages for data analysis and machine learning based on rough sets, are the primary reasons behind our choice of the rough sets approach to machine learning in the experiments reported in this paper. Second, over the last 2 years, good cooperation has been established between the Intelligent Computers Laboratory at Wayne State University and Reduct Systems, Inc., of Regina, Canada. As a result of this cooperation, we have developed a good understanding of the methodological aspects of using learning systems based on rough sets in engineering. Also, these systems have proven to be reliable and user friendly, and they have special features that were developed to address our specific engineering needs, for example, a component for automatic knowledge verification.

However, there are several other classes of experimental and commercial learning systems that could be used in our project. For example, the learning algorithm ID3 proposed by Quinlan (1986) has been implemented in SuperExpert, a commercial learning system developed by Intelligent Terminals (1986). Also, a class of learning algorithms AQ proposed by Michalski (Michalski & Chilausky, 1980) has been implemented in various experimental learning systems developed at the center for Artificial Intelligence Research at George Mason University, Fairfax, Virginia. The possibility is considered that these systems, in particular, could be used for our experiments. Therefore, at present it would be premature to claim that the selected learning system has an absolute superiority to the other systems available. However, the performance of various learning systems on the transportation engineering data should be investigated in the future.

Our research was planned to address existing research needs in a specific subarea of civil engineering. However, this subarea was also selected for general reasons. From the methodological point of view, the control of traffic in a rail corridor is similar to other types of traffic control, e.g., air, road, or sea. Also, rail traffic control is a specific form of process control, and therefore, our conclusions can, at least partially, be generalized for other forms of process control, including the control of manufacturing processes. Thus, we hoped that our research would produce results useful for both transportation engineers and engineers in other specialties interested in the control of complex processes and in the use of learning systems to acquire knowledge about these processes.

METHODOLOGY OF KNOWLEDGE ACQUISITION

This research is one of the first attempts to use machine learning in transportation engineering, and therefore, no domain-specific methodological experience was available that could be used to design the knowledge acquisition process. For this reason, significant attention was paid to studying known processes of automated knowledge acquisition in civil engineering, in order to develop a proper methodology of knowledge acquisition for our specific area of rail corridor traffic control. Finally, our methodology was partially based on that developed for automated knowledge acquisition in structural engineering (Arciszewski et al., 1987; Mustafa & Arciszewski, 1992; Reich & Fenves, 1992) and partially on the general engineering methodology of automated knowledge acquisition proposed by (Arciszewski and Mustafa (1989). Also, recent results were utilized on the development of learning engineering at the Center for Artificial Intelligence Research at George Mason University. Learning engineering is a new subarea of knowledge engineering, which deals with the methodological aspects of using learning systems in knowledge acquisition in science and technology. It encompasses the evaluation and selection of learning systems, the methodology of knowledge acquisition, and the verification of knowledge. In our project, the empirical error rates developed as a part of the method for the performance-based evaluation of learning systems (Arciszewski et al., 1992) were used for knowledge verification. As a result of our studies, the following eight-stage knowledge acquisition process was designed:

1. learning about the domain
2. methodological studies

3. development of the representation space
4. development of mathematical models
5. implementation of models
6. preparation of examples
7. learning decision rules
8. knowledge verification

In the first stage, learning about the domain, the available literature on the control of traffic in an urban rail corridor was analyzed, and a list was prepared of the most appropriate sources of knowledge, including books, papers, and research reports that were found potentially useful. All this material was discussed by the transportation engineers on our team, and a final collection of relevant publications was identified. Next, these publications were studied to learn the state-of-the-art in the area of interest, to identify available formal models for the control of rail traffic, and to find an initial collection of attributes to describe our problem.

In the second stage, methodological studies, an analysis of similar feasibility studies regarding the use of machine learning in civil engineering was conducted, including studies in the areas of architectural building design (Gero et al., 1989), design of tall buildings (Arciszewski & Ziarko, 1988), design of bridges (Reich & Fenves, 1992), and construction safety (Arciszewski & Usmen, 1993). Several computer scientists who work on machine learning and its engineering applications were contacted for their suggestions and input. The most important response was that the representation space should be particularly carefully prepared and that nominal binary attributes should to used, if possible, to improve the performance of a learning system on actual engineering examples. Both recommendations were followed, and the development of knowledge representation space even became a separate stage in the knowledge acquisition process. This stage resulted in the design of our knowledge acquisition process and provided a methodological framework for our project.

As expected, the third stage, development of representation space, was particularly difficult. Because of its importance, it is discussed separately in some detail below. Similarly, stages three through five were difficult and time consuming, and are discussed jointly in the section on preparation of examples.

In the seventh stage, learning decision rules, all the examples prepared were utilized to produce decision rules using DataLogic, which is briefly described in the section on knowledge acquisition. In the eighth stage, knowledge verification, the decision rules produced were verified using empirical error rates assumed in accordance with the work of Arciszewski et al. (1992). These error rates and the results of the knowledge verification are given in the section on knowledge verification.

Our knowledge acquisition process was adequate for the problem of learning about rail corridor traffic control. It is relatively simple, but at the same time

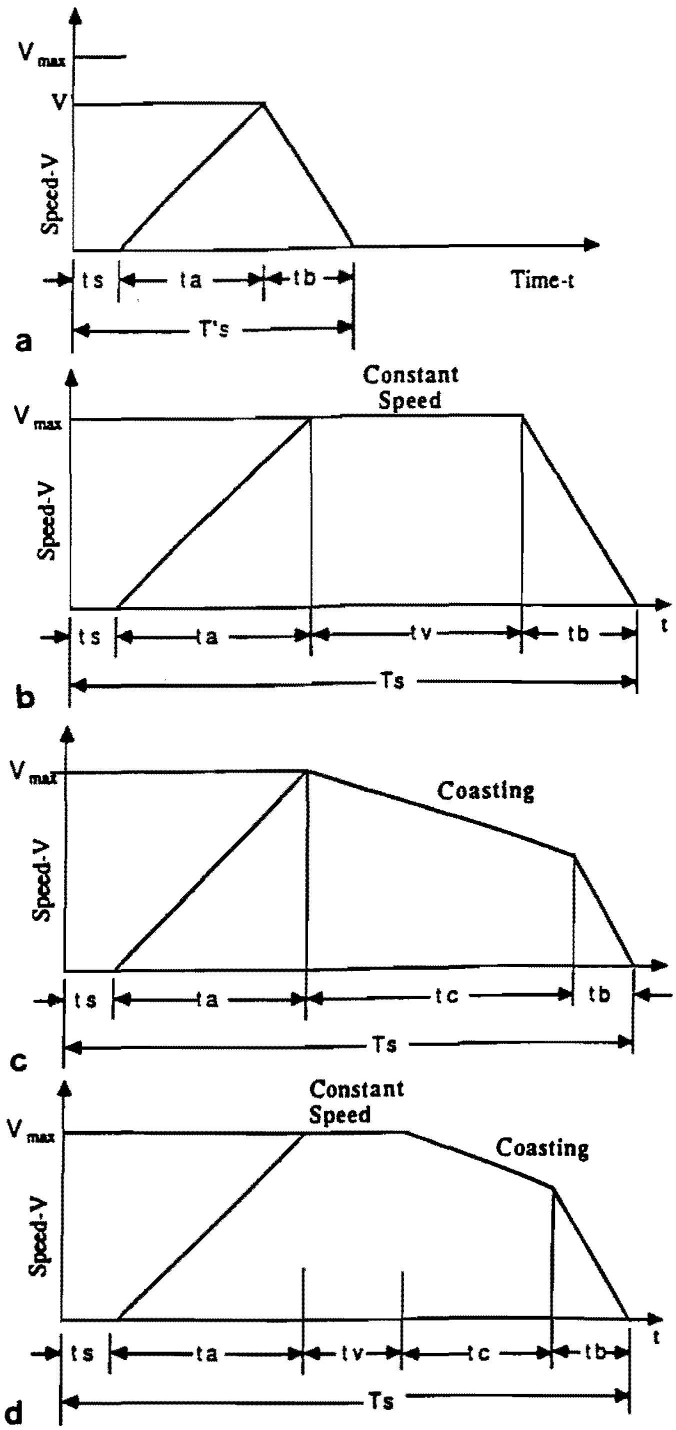

Figure 1. Interstation travel regime. (a) Regime A; (b) Regime B; (c) Regime C; (d) Regime D. (Vuchic, 1980)

sufficient to produce the expected results, and therefore, it can be recommended for other knowledge acquisition projects in the area of transportation engineering.

DEVELOPMENT OF REPRESENTATION SPACE

The learning about the domain stage in our knowledge acquisition process resulted in a good understanding of traffic control along an urban rail corridor with closely spaced stations, in the context of both decision making and machine learning. It has been concluded that this traffic can be considered as a sequence of decision-making stages. When a train on an urban corridor connecting two terminal points with a large number of intermediate stations is considered, it can follow various regimes of motion and stopping, which identify individual driving scenarios. By a scenario, we mean a unique combination of decisions (decision attributes and their nominal values) that are taken into account by the traffic controller. Execution of individual scenarios may result in different values of performance attributes to describe driving scenarios from the performance point of view, namely, travel time, energy consumption, and comfort. In this case, decision attributes, which are controlled by the urban rail corridor operator, can be considered as independent attributes. The performance attributes are only indirectly controlled by the operator through decision attributes, and they can be considered as dependent attributes. The entire problem of control of an urban rail corridor can then be described by two classes of attributes: independent and dependent. The development of the representation space for this problem must then involve the identification of all attributes, both independent and dependent, and the determination of their values, preferably nominal.

Regime of Motion

Any urban rail corridor can be divided into a number of segments. The operator selects suitable regimes of motion for individual segments from four basic regimes, called A, B, C, and D, which are as follows (Vuchic, 1980):

Regime A: The interstation spacing is less than the critical spacing; critical spacing is the minimum spacing between stations needed for the train to attain its maximum speed (Figure 1a).

Regime B: The interstation spacing is longer than the critical spacing. The train maintains a sustained level of maximum speed before deceleration is initiated for the next stop (Figure 1b).

Regime C: The interstation spacing is longer than the critical spacing. However, as an energy conservation measure, the train starts coasting (decelerating at a very slow rate) immediately upon reaching its maximum speed, and continues to travel at coasting speed until deceleration is initiated as the train approaches the next station (Figure 1c).

Regime D: This regime is a combination of regimes B and C. It allows the train to travel at its maximum speed for some time and to coast between two stops. Within regime D, an infinite number of combinations is possible, depending on the instant when coasting is initiated. The limiting cases are regime B if coasting begins immediately prior to braking and regime C if coasting is initiated immediately upon the attainment of maximum speed.

Energy Consumption

Studies of Hamburg rail systems by empirical and computer simulation techniques have demonstrated the importance of different driving regimes for energy consumption (Mies, 1969). The trade-off between energy consumption and travel time was developed from time-speed-energy consumption data. The results were used in this study to develop surrogate measures of energy consumption for varying travel times in the form of an empirical relationship between time and energy consumption. Although this relationship does not explicitly consider different regimes of motion described above, lower energy consumption resulting from longer coasting and consequent longer travel times are incorporated in the above relationship (Mies, 1974).

A total of four models were developed for estimation of energy consumption using the Hamburg data (Vuchic, 1980). These models were a simple model, a polynomial model, a logarithmic model, and an exponential model. The following exponential model was used for the study:

$$Y = 1322.5 \times 10 \wedge (-1.0097e^{-2X}) \qquad R^2 = .983$$

where X is travel time surrogate and Y is energy consumption surrogate.

Passenger Comfort Levels

Every change in acceleration/deceleration phase is associated with a level of discomfort for the passenger. A change in the rate of acceleration/deceleration (second derivative of speed with respect to time) is commonly termed a "jerk." Therefore, it has been assumed that the level of discomfort experienced by a passenger is measured by the number of jerks during a given pass of the train along the entire corridor. Ideally, not only the frequency but also the respective magnitudes of jerks should be considered. However, magnitudes were considered too complex to quantify for the purpose of this study.

For a typical interstation travel, two instances of jerks will be experienced during the acceleration phase, two during the deceleration phase, and one during the beginning of the coasting operation. Further, for every skip-stop operation a total of four instances of jerk can be "saved," resulting from the elimination of deceleration and acceleration operation as the train approaches and leaves the station in question.

Regime Simulation Model

When controlling a given urban corridor, the operator usually selects driving scenarios to optimize one of the three performance attributes, e.g., to minimize travel time or energy consumption, or to maximize comfort. The values of the performance attributes for individual driving scenarios are usually calculated in advance, using optimization models prepared to deal separately with travel time, energy consumption, and comfort. A simulation model called REGIME developed for our project, is briefly discussed in the following section (Khasnabis et al., 1992).

The general model of the control of an urban rail corridor with any number of stations and a large number of driving scenarios would be complex and would require the use of many attributes to describe it. This would, consequently, lead to a large representation space, and learning over this space would require a large number of examples. Therefore, the development of the representation space and examples would be highly complex. Since our objective was to conduct a feasibility study only, we decided to consider a hypothetical urban rail corridor and to learn traffic control decision rules for this corridor. However, our hypothetical rail corridor was assumed as a realistic example of an urban rail corridor of average complexity.

An urban rail corridor of fifty intermediate station-spaces (sections) was assumed. The corridor was divided into five segments (Figure 2). Segments one and five are end segments, each consists of four station spaces. Segments two and four are the two intermediate segments, each consisting of twelve station spaces. Segment three is the central segment and contains eighteen station spaces. Each spacing was assumed to be 2000 ft for a total corridor length of 100,000 ft. The urban rail corridor analyzed is thus symmetrical, with segments one and two being mirror images of segments five and four respectively, and segment three being the central portion of the corridor.

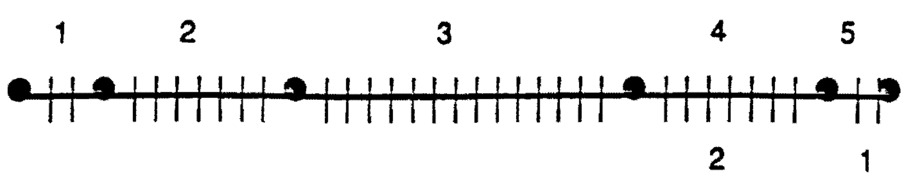

Figure 2. Schematic of the study corridor consisting of five segments. Segment 1 and segment 5 are symmetrical. Segment 2 and segment 4 are symmetrical. Segment 1 and 5 have four interstation spacings at 2000 ft each. Segment 2 and 4 have twelve interstation spacings at 2000 ft each. Segment 3 has eighteen interstation spacings at 2000 ft each. Total length is 100,000 ft with 50 station spacings.

It was assumed that the decisions made for the first and fifth segments would be identical. Similarly, it was assumed that the second and fourth segments were controlled by identical decisions. Therefore, the rail corridor was completely described when decisions for the first, second, and third segments were known (values of decision attributes or independent attributes). For each of these three segments, five decisions concerning train operations are to be made. Thus, the entire process of train control can be considered as a sequence of fifteen decisions. The nature of the individual decisions was determined using domain knowledge, and they were designed as binary decisions to improve the performance of the learning system on examples regarding our problem, as suggested by one of our machine learning consultants. These binary decisions require YES or NO answers, and the independent attributes representing these decisions are given in Table 1. Thus the representation space includes fifteen binary nominal independent attributes and three nominal dependent attributes. The total number of possible scenarios is then (2)∧15, 32,768.

PREPARATION OF EXAMPLES

All examples were prepared for a hypothetical urban rail corridor, as described in the preceding section. For each driving scenario the values of the performance attributes were calculated using REGIME, a computer program developed for the purposes of our research.

REGIME has three basic components, which separately produce estimates of performance attributes for a given train driving scenario along an assumed corridor.

Table 1. Knowledge representation

Segment	Attribute	Attribute values		
1 and 5	constant speed	yes		no
	constant speed and coasting	yes		no
	one stop skipped	yes		no
	two stops skipped	yes		no
2 and 4	constant speed	yes		no
	coasting	yes		no
	constant speed and coasting	yes		no
	one stop skipped	yes		no
	two stops skipped	yes		no
3	constant speed	yes		no
	constant speed and coasting	yes		no
	one stop skipped	yes		no
	two stops skipped	yes		no
Passenger comfort		high	medium	low
Travel time		high	medium	low
Energy consumption		high	medium	low

For each scenario and performance attribute, REGIME conducts the analysis of various train operations to produce an estimate of the performance attribute considered. The component for the analysis of travel time was developed for individual regimes using the time and travel speed analysis algorithms developed by Vuchic (1980). The analysis of energy consumption was based on an exponential model of the relationship between travel time and energy consumption, which was developed as a part of earlier research (Mies 1969, 1974). It was assumed that comfort can be measured by the number of jerks.

REGIME was used to analyze 102 representative scenarios, for which values of all three performance attributes were separately produced. These scenarios were selected using the domain experience of the transportation engineers on our research team. The number of scenarios was very small when compared with the size of the representation space: only 0.3% of the representation space was covered by the scenarios. However, the scenarios were "balanced," in that they were carefully selected to cover the entire representation space uniformly (Arciszewski & Mustafa, 1989). Balanced examples were chosen in our study because their use significantly improves the performance of a learning system more so than use of randomly selected examples (Mustafa & Arciszewski, 1992).

REGIME produced estimates of the performance attributes in the following ranges:

travel time (seconds)	2057–3709
energy consumption (surrogate)	64–129 units
passenger comfort (number of jerks)	72–250

Two levels of categorization were used, one consisted of subranges 1–3 of equal length, and the second consisted of subranges 1–7 of equal length. In this way, two collections, each with 102 examples, were produced for each performance (dependent) attribute. The first collection was with three categories of the dependent attribute and the second with seven categories.

In our study, two different numbers of categories of the dependent attribute were considered, to determine the sensitivity of the performance of a learning system to the number of dependent attribute categories. For rail corridor control, the larger number of dependent attribute categories has many advantages, including more precise decision making and decision rules that may be more specific. However, we were aware that increasing the number of dependent attribute categories usually worsens the performance of a learning system. The trade-off between the number of dependent attribute categories and the domain usefulness of the decision rules produced using these categories is important in engineering, and therefore any result improving this trade-off can be useful.

For our experiments, six collections of 102 examples each were prepared, two for each of the three performance attributes considered. In total, 612 examples were developed.

KNOWLEDGE ACQUISITION

Learning was conducted using the commercial system DataLogic, which was developed by Reduct Systems, Inc., of Regina, Canada. DataLogic is a general-purpose software package for automated knowledge acquisition and for building classification expert systems. The learning component of the system is based on a learning methodology derived from the mathematical theory of rough sets (REDUCT Systems, 1991; Pawlak, 1992). The particular algorithm implemented in DataLogic performs the major processing stages. It accepts a relational table of training data. The user is required to mark a selected subset of attributes as conditions, and exactly one attribute is marked as a decision. The decision attribute represents the classification the system is trying to learn. The generation of rules involves the following steps:

1. Formation of higher order attributes to represent the original information. Essentially, the precision of original data is reduced by replacing the individual values with range symbols after dividing each attribute domain into a predefined set of ranges.
2. Classification of the original data into identified classes based on the range attributes.
3. Elimination of redundant attributes, which are the attributes that can be eliminated without affecting the quality (accuracy) of classification. The rough-set concept of reduct (Pawlak, 1992) is used at this stage.
4. Reclassification of the original data based on the reduced collection of attributes.
5. Elimination of redundant attribute values from the collection of reclassified data obtained in step 4. At this stage, the concept of "value reduct" (Pawlak, 1992) is used. After step 5, a set of decision rules is produced. The details of the rule extraction procedure can be found in the work by Pawlak (1992).

Six separate automated knowledge acquisition processes were conducted, each for a different collection of examples, as described in the previous section, on preparation of examples. For individual processes, different collections of decision rules were obtained. Numbers of these decision rules in all collections are given in Table 2.

It has been observed that in all cases the use of a larger number of dependent attribute categories results in a larger number of decision rules, as expected. In our experiments, the average relative difference calculated with respect to the number of decision rules for a collection with a smaller number of dependent attribute categories is approximately 33%. However, the complexity of the decision rules, in terms of the number of independent attributes used in these decisions, is surprisingly comparable, although different combinations of attributes are used. For example, when the dependent (performance) attribute "travel time" was considered, the

Table 2. Numbers of decision rules for individual learning cases

Performance attribute	Automated knowledge acquisition process	No. of dependent attribute categories	No. of decision rules	Relative change in the no. of decision rules (%)
Travel time	1	3	37	38
	2	7	51	
Energy consumption	1	3	34	15
	2	7	39	
Passenger comfort	1	3	38	47
	2	7	56	

following two corresponding decision rules were obtained for the cases with three and seven values of the performance attribute, respectively. When

$$S14 = 1 \qquad S24 = 1 \qquad S32 = 0$$

then

$$\text{travel time} = 2056\text{–}2607$$

or to achieve travel time between 2056 s and 2067 the following conditions must be fulfilled.

- There *must* be one skip-stop (S14) in the first and fifth segment.
- There *must* be one skip-stop (S24) in the second and fourth segment.
- There *must not* be any coasting (S32) in the third or middle segment.

And when

$$S25 = 1 \qquad S21 = 1 \qquad S11 = 1$$

then

$$\text{travel time} = 2056\text{–}2293$$

or to achieve travel time between 2056 s and 2293s, the following conditions must be fulfilled:

- There *must* be two skip-stops (S25) in the second and fourth segment.
- There *must* be constant speed (S21) in the second and fourth segment.
- There *must* be constant speed (S11) in the first and fifth segment.

All the collections of decision rules obtained were analyzed from the domain point of view. They are clear in operational terms, but understanding and explaining them in the context of state-of-the-art urban rail control is difficult. A review of the decision rules indicates that, as individual entities, they are logical and rational. However, the exact interpretation of these rules collectively is a matter of further research.

KNOWLEDGE VERIFICATION

In our project, we assumed that knowledge in the form of decision rules can be formally verified using empirical error rates that were initially developed for the evaluation of the performance of learning systems (Arciszewski et al., 1992). Therefore, the knowledge produced by the learning experiments described in the preceding section was verified, using the overall empirical error rate determined by multisampling conducted using the "leave-one-out" method (Weiss & Kulikowski, 1991; Arciszewski et al., 1992). This error rate is considered the most significant empirical error rate that can be used to evaluate the performance of a learning system and to verify the knowledge automatically produced, as demonstrated by Arciszewski and Dybala (1992).

The overall empirical error rate with multisampling using the "tenfold" method was also considered for use in knowledge verification, but was finally rejected. This decision was based on our observation (Arciszewski & Dybala, 1992) that when the number of examples is relatively small, as in our case, the overall error rate determined using the tenfold method is usually not significant, because it may undergo large changes when different groups of examples are selected.

The overall empirical error rate is defined as

$$E_{ov} = \text{number of errors/number of tests}$$

where "error" is a misclassification of a testing example and "number of tests" is the number of classification tests.

In the "leave-one-out" resampling method (Weiss & Kulikowski, 1991), the empirical error rate is calculated as an average error rate for n repetitions (n is the number of examples). In each repetition a different example is removed from the collection of n examples, and the remaining $n - 1$ examples are used to produce decision rules, which are utilized next to predict the example that was removed.

Automatic knowledge verification was conducted for all six collections of decision rules produced in our experiments. The calculations were performed using a special DataLogic component for knowledge verification. All values of the overall empirical error rates (the leave-one-out method) are shown in Table 3.

A significant difference in the performance of the learning system, and of the quality of decision rules measured by this performance, can be observed between decision rules produced for the case of three values versus the case of seven values of the dependent attribute. The average relative difference, calculated with respect

Table 3. Overall error rates for individual learning cases

Performance attribute	Automated knowledge acquisition process	No. of dependent attribute categories	Overall empirical error rate (%)	Relative difference in error rates (%)
Travel time	1	3	14	43
	2	7	20	
Energy consumption	1	3	12	40
	2	7	20	
Passenger comfort	1	3	9	67
	2	7	15	

to results obtained for the cases with three dependent attribute categories, is approximately 50%. This result is not surprising, and it clearly demonstrates the well-known heuristic that the performance of a learning system and the quality of decision rules produced by it rapidly deteriorate when the number of categories of the dependent attribute is significantly increased.

In our experiments, a relatively small number of examples (0.3% coverage of the representation space) was used, and therefore, the overall empirical rate obtained was in the range 9–20%, which seems relatively high. However, this result should be compared with the performance of human experts, who usually make decisions regarding complex engineering problems with accuracies of 70–80%, i.e., with an overall rate of 20–30%. In the research reported, this comparison was not conducted, but it is planned for the future.

CONCLUSIONS

The feasibility study reported in this paper was the first of its kind in transportation engineering. It took approximately 12 months to complete the project, with a graduate research assistant working 50% of his time during the academic year, and the remaining members of the team being involved on an irregular basis, as needed. The project successfully demonstrates the feasibility of using machine learning in knowledge acquisition process and to the methodological aspects of building representation space. The preparation of examples required conducting extensive domain studies and the development of mathematical models for the analysis and optimization of travel time, energy consumption, and passenger comfort, which were used in REGIME, a computer program prepared as a part of our research.

REFERENCES

Arciszewski, T., and T. Dybala. 1992. *Evaluation of learning systems: a method and experimental results*. Fairfax, Va.: Center for Artificial Intelligence, George Mason University.

Arciszewski, T., T. Dybala, and T. Wnek. 1992. A method for evaluation of learning systems. Heuristics, special issue on machine learning in knowledge acquisition. *J. Knowledge Eng.*

Arciszewski, T., and M. Mustafa. 1989. Inductive learning process: the user's perspective. In *Machine Learning*, ed. R. Forsyth. London: Chapman and Hall.

Arciszewski, T., M. Mustafa, and W. Ziarko. 1987. A methodology of design knowledge acquisition for use in learning expert systems. *Int. J. Man Mach. Stud.*, no. 27.

Arciszewski, T., and M. Usmen. 1993. Applications of machine learning to construction safety. Paper presented at the International Conference on Management of Information Technology for Construction, Singapore.

Arciszewski, T., and W. Ziarko. 1988. Adaptive expert system for preliminary design of wind bracings in skeleton structures. In *Second century of skyscrapers*. Princeton, N.J.: D. Van Nostrand.

Gero, J., C. H. Mackenzie, and S. McLaughlin. 1989. Learning from optimal solutions to design problems, ed. B. Topping. Paper presented at the NATO ASI Conference on Optimization and Decision Support Tools in Civil Engineering, Edinburgh.

Khasnabis, S., T. Arciszewski, and S. Hoda. 1992. Learning rules for driving scenarios for an urban rail corridor with closely spaced stations. Proceedings of the ASCE Conference on Computing in Civil Engineering, eds. B. J. Goodno and J. R. Wright, Dallas, Tex.

Michalski, R. S., and R. L. Chilausky. 1980. Learning by being told and learning from examples: an experimental comparison of the two methods of knowledge acquisition in the context of developing an expert system for soybean disease diagnosis. *Policy Anal. Inform. Syst.* 4:125–160.

Mies, A. 1969. *Fahrdynamische Grundueberlegungen fuer die Automatisierung von Schnellbahnen*. Bielefeld, Germany: Vevhkehr und Technic, Erich Schmidt Verlag.

Mies, A. 1974. *Underground railway automation of the Hamburger Hochbahn AG*. UITP revue no. 1, 45–48. Brussels: UITP.

Mobility 2000 Working Group. 1990. Intelligent vehicle highway systems: operational benefits.

Modesitt, K. L. 1992. Basic principles and techniques in knowledge acquisition. In *Knowledge acquisition in civil engineering*, eds. T. Arciszewski and L. Rossman. New York: American Society of Civil Engineers.

Mustafa, M., and T. Arciszewski. 1992. Inductive learning of wind bracing design for tall buildings. In *Knowledge acquisition in civil engineering*, eds. T. Arciszewski and L. Rossman. New York: American Society of Civil Engineers.

Pawlak, Z. 1992. *Rough sets: theoretical aspects of reasoning about data*. Norwell, Mass.: Kluwer Academic.

Quinlan, R. 1986. Induction of decision trees. *J. Mach. Learning* 1(1):81–106.

Reich, Y., and S. Fenves. 1992. Automated design knowledge acquisition by concept formation. In *Knowledge acquisition in civil engineering*, eds. T. Arciszewski and L. Rossman. New York: American Society of Civil Engineers.

Reduct Systems. 1991. *DataLogic—user's manual*. Regina, Canada.

Slowinski, R. 1992. *Intelligent decision support*. Norwell, Mass.: Kluwer Academic.

Softsync, Inc., 1986. Intelligent terminals. In *SuperExpert, manual*. New York.

Underwood, S. E., K. Chen, and K. Ervin. 1989. *The future of intelligent vehicles highway systems: a delphi forecast of markets and sociotechnial determinants*. IVHS report TR2-89. Ann Arbor: University of Michigan.

Vuchic, V. R. 1980. *Urban transportation: systems and technology*. Englewood, Cliffs, N.J.: Prentice-Hall.

Weiss, S. M., and C. A. Kulikowski. 1991. *Computers that learn*. San Mateo, Calif.: Morgan Kaufmann.