

Real-Time Data Fusion for Arterial Street Incident Detection Using Neural Networks

JOHN N. IVAN, JOSEPH L. SCHOFER, FRANK S. KOPPELMAN,
AND LINA L. E. MASSONE

This research contributes to the development of an automatic incident detection system for detecting traffic-delaying events on arterial street networks for an Advanced Traveler Information System demonstration called ADVANCE. Data describing current traffic conditions will be gathered in real-time from two distinct sources: inductive loop detectors and specially equipped vehicles that measure and report their travel times on roadway links. Two approaches are considered for data fusion, the combination of information from these sources to produce a single decision about the presence or absence of incidents on each link. In the integrated fusion approach, observed traffic data are combined directly using a neural network. In the algorithm output fusion approach, separate incident-detection algorithms individually preprocess data from each source, reporting outputs that are combined using a neural network. Data for calibrating these system components were generated using computer simulation. The algorithm output fusion network performed better than the other approach, detecting over 80 percent of the incidents with almost no false alarms. Fusing algorithm outputs using neural networks was thus found to improve the capability provided by separate source incident detection algorithms operating alone. The importance of validating these results through calibration and testing with field data, as well as improving performance through introduction of an additional data source is discussed.

Highway facilities are designed to operate acceptably within some range of demand. When an incident occurs, such as a traffic accident, a load spill, or a vehicle breakdown in the roadway, there is a sudden, temporary decrease in the capacity of a particular section of the facility. When demand exceeds this temporarily reduced capacity, queues, delays or perhaps more accidents result, as well as increased difficulty in clearing the scene (1).

Identifying incidents quickly is important for reducing their impacts. Researchers over the years have developed Automatic Incident Detection (AID) systems for freeways that monitor traffic flow information from a highway facility and automatically detect such incidents (2) to prevent secondary accidents from occurring and to dispatch emergency or cleanup crews promptly. A few researchers have studied incident detection on arterial streets (3,4,5), using techniques similar to those used on freeways. However, the arterial street environment is much more challenging because traffic flow discontinuities are introduced by traffic signals, traffic entering and leaving side streets and driveways, and variation in signal timing and geometric characteristics. Therefore, the effect of the same type of incident varies for different sets of arterial conditions.

Recently, there has been a great deal of interest in increasing the efficiency of existing highways through the deployment of Advanced Traveler Information Systems (ATIS), a type of Intelligent Transportation System (ITS). An ATIS provides current traffic information to system users to help them reduce their travel times. The Illinois University Transportation Research Consortium (IUTRC), Motorola, and the Illinois Department of Transportation (IDOT), with funding from the Federal Highway Administration (FHWA), are preparing to launch a demonstration of such a system in the Chicago suburbs, called ADVANCE, or Advanced Driver and Vehicle Advisory Navigation Concept (6). ADVANCE will provide between 3,000 and 5,000 demonstration participants with shortest path routings to specific destinations using up-to-the-minute travel times on arterial and freeway links in an approximately 600 sq km service area. Advising drivers of traffic operational problems on the highway network will be an important function of the demonstration, so an AID system will be important for enhancing the value of information provided by the system and promoting constructive driver response (7). Detailed information about current traffic conditions and causes will be particularly valuable for predicting travel times.

METHODOLOGY

Data Sources

AID systems described in the literature have operated almost exclusively using data from fixed detectors, sensor systems that measure traffic characteristics at a fixed location, such as inductive loop detectors (ILDs) or video cameras (2). Fixed detectors measure the following traffic flow quantities:

1. Volume, the arrival rate of vehicles passing the detector during the measurement period;
2. Occupancy, the percentage of time that the space above the detector is occupied by a vehicle during the measurement period; and
3. Speed, the rate of motion of a vehicle as it passes the detector (the detector may provide individual vehicle speeds or average speed over the measurement period).

Traffic signals on primary arterial routes (on which traffic flow is most vulnerable to unusual congestion) are often connected in closed loop systems, which coordinate signal timings and facilitate collection of data in real-time over data communication lines. Not all of the roadway links in the ADVANCE street network will be instrumented with these systems, so additional data sources are required to cover other network sections.

J. N. Ivan, University of Connecticut, Transportation Institute U-37, Storrs, Conn. 06269-2037. J. L. Schofer and F. S. Koppelman, Northwestern University Transportation Center, Evanston Ill. 60208. L. L. E. Massone, Northwestern University, Departments of Electrical Engineering and Computer Science, Department of Biomedical Engineering, Evanston, Ill. 60208

One potential alternative source is observed travel times collected in real-time from probe vehicles traveling the street network, a common feature of many ITS implementations, including ADVANCE. Vehicles participating in ADVANCE will automatically report observed link travel times to a Traffic Information Center (TIC). Probe vehicles can help locate congestion on any roadway segment in the area served by the vehicle communications medium, without the spatial limitations of fixed detectors, although network coverage is limited by the market penetration rate.

ADVANCE will collect data from both of these sources, which are considerably different from each other both qualitatively and temporally. It is thus desirable to use separate procedures, or incident detection algorithms, for each data source to determine whether or not there is an incident (or the likelihood of there being an incident) on each link. A data fusion process would then solve the more clearly defined task of combining the incident decisions made by the two procedures.

Data Fusion Approaches

The focus of this research is the development of this data fusion process. These data sources are inherently imperfect and not entirely reliable, so this procedure must be able to:

1. Identify incident conditions under a variety of input data patterns,
2. Integrate inferences from input data with varying degrees of certainty; and
3. Account for complex relationships among input sources.

A number of information processing techniques have proven useful for decision-making and combining uncertain information in a variety of contexts. Following is an evaluation of the potential effectiveness of several such techniques for solving the data fusion problem posed here.

Decision support systems interpret surveillance information and recommend a course of action for the system operator, who must then make a decision. Prosser and Ritchie (8) describe such a system for incident management. This system really filters available information, conserving the operator's attention for confirming computer generated results. While it might be necessary for an operator to monitor incident detection system operation to avoid broadcasting false alarms, particularly in the early stages of implementation, it is more desirable for the ADVANCE incident detection system to operate without requiring regular operator response.

The best score approach (9) simply chooses the information source that is considered most valid (according to a predetermined quality score and an aging or decay rate) and uses it alone. This procedure may be better described as a "winner-take-all" strategy rather than data fusion. Discarding other information ignores possible interactions among the various data sources that contribute to system performance.

Virtual sampling (10) regards several estimates of an unknown quantity as unique random samples of observations drawn from a given population whose mean value is the unknown quantity. Estimates with low standard errors are considered to represent larger samples. For combining data sources that represent independent observations of the same phenomenon, this approach seems quite attractive, since it places more emphasis on values that are more certain and requires no weights to be calibrated. However, it does require that standard errors be provided with each estimate.

Artificial Neural Systems (ANS), also known as neural networks, are information processing structures that attempt to replicate the process of learning and decision-making observed in the operation of the human brain (11,12). The nature of neural networks makes them appropriate for solving many complex problems that have proven to be quite cumbersome for conventional processing systems. For example, problems in which the precise interrelationships among elements are not well understood, such as continuous speech processing and pattern recognition, are good applications for ANS. An ANS is best implemented as a partner to a traditional system, with the traditional system (for example, our incident detection algorithms) performing precise, specific computations and the ANS analyzing less precisely defined tasks (such as data fusion).

Approach Concepts

Two data fusion approaches using neural networks are considered here.

Algorithm Output Fusion

In this approach, depicted in Figure 1, two algorithms, each uniquely developed for one of the two data sources (fixed detectors and probe vehicles), determine the likelihood that an incident is occurring at particular locations on the street network. A separate data fusion process using a neural network then combines the output from these algorithms.

Integrated Fusion

Here the functions of the single source incident detection algorithms and the fusion process are combined in a single neural network. In the first approach, the algorithms effectively censor the unprocessed input data, translating them into a single output value, preventing unprocessed input data from one source from helping to interpret data from the other source. This network will read input directly from the data sources, then fuse data and detect incidents simultaneously.

NETWORK TRAINING

Training Data

Humans learn by repeatedly observing the outcomes of their responses to external stimuli. In the same way, network training in

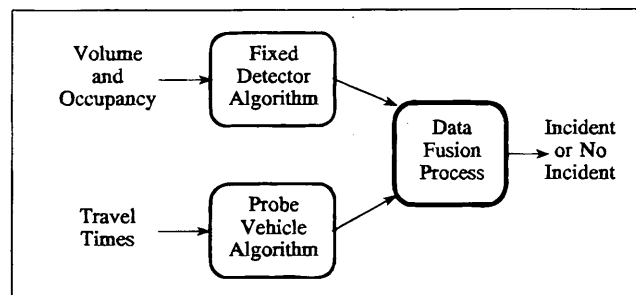


FIGURE 1 Algorithm output fusion concept.

this application requires a set of data that represent the variety of traffic conditions under which the incident detection system must operate, along with the outcomes it should return.

Most incident detection systems are calibrated for specific road sections: a different set of parameter values—for example tolerance thresholds—is used for each pair of detectors, or road section (2). For an arterial street network with many road segments, such a calibration would be exceedingly tedious. Instead, for this incident detection system, it was desired to calibrate more general algorithms that would apply to any of the arterial segments in the service area.

To do this, the training data must incorporate the variation that arises naturally in traffic flow on urban arterial streets. This implies an enormous number of combinations of street, traffic, driver, and incident patterns, of which the data set must be a suitable sample. Much of this variability is derived from human behavior (driving patterns, the occurrence of incidents), and real traffic conditions are the best source for unbiased observation of these phenomena. However, real-world traffic data involving incidents are considerably difficult to obtain (particularly for the sources involved here), so the training data used here were generated through traffic simulation using INTRAS, a microscopic freeway corridor traffic simulation model (13).

The arterial street network simulated for data collection is a representation of an approximately 5-km section of major arterial streets in the ADVANCE network, as depicted in Figure 2. This road section has 39 loop (fixed) detectors at eight intersections that are located to collect data for signal controllers. The simulation model includes these detectors along with additional detectors such that detector stations are located at each eastbound intersection approach. Signal timing was varied at several intersections to study the effect of incidents under different congestion (volume to capacity ratio) conditions. Similarly, incidents were placed at different locations on links to study the effect of the distance from the signal on traffic operation.

More than 100 simulation runs were performed to generate data for a variety of incident and corresponding nonincident conditions. A data aggregation interval of 7 min was selected so that no cycles would run between two intervals (all signals have cycle lengths of 140 sec); flow variations through the cycle thus do not taint the traffic measurements, but the interval is short enough to permit timely traffic condition updates. For each incident simulation, a number of nonincident simulations with identical control variables were also performed. Incidents were simulated on six different links, at three or four locations on each link, for durations of from 5 to 10 aggregation intervals, and with up to three different signal timing patterns at selected signals. All of the incidents were simulated in the eastbound direction, so only data from the eastbound links are used in the analysis.

Training data were prepared by extracting aggregated occupancy, volume, and travel time reports from INTRAS output corresponding to each simulated incident. For the Algorithm Output Fusion Network, input vectors were generated by processing traffic data with the calibrated single source incident detection algorithms (14). Each input vector corresponds to the conditions on one link during a particular incident simulation time interval, and consists of the fixed detector and probe vehicle algorithm scores scaled to keep their values between -1.0 and $+1.0$ (negative values indicate no incident, positive values indicate an incident) and the target output equal to 1.0 if there was an incident on the link during the time interval and 0.0 if not.

A similar process was used to organize the simulated traffic surveillance data into training vectors for the integrated fusion approach, but the following input values replace the algorithm scores on each vector:

1. Volume ratio, equal to the traffic flow at the detector station during the time interval divided by the average total traffic flow at the station under nonincident conditions;
2. Occupancy ratio, equal to the average traffic occupancy measured at the station on the analysis link during the time interval divided by the average traffic occupancy for the same station under nonincident conditions; and
3. Travel Time Ratio, equal to the average of the travel times observed on the analysis link during the time interval divided by the average travel time for that time period on the link under nonincident conditions.

All of these values were also scaled to keep their values between 0 and 1.

Training Procedure

Both data fusion approaches were developed with feed forward networks trained using error back propagation (12). Training prepares a network for application and involves presenting a series of input arrays, or vectors, to the network one at a time along with their corresponding target output values. The network adjusts connection weights over a series of many epochs, or iterations, so that it can reproduce the desired output values.

The network structure for the two approaches are depicted in Figures 3 and 4. Both use a single hidden layer of five units and a single output unit. All hidden layer units and the output unit add a bias threshold to their net inputs. For each input vector presentation, the output value is calculated using Equation 1

$$y = f\left(v_0 + \sum_{j=1}^m v_j f\left(w_{0j} + \sum_{i=1}^n w_{ij} x_i\right)\right) \quad (1)$$

where

- y = resulting network output;
- f = activation function;
- v_0 = bias threshold on the output unit;
- m = number of hidden units;
- v_j = weight on the connection from hidden unit j to output unit;
- w_{0j} = bias threshold on the hidden unit j ;
- n = number of input units;
- x_i = input signal from input i ; and
- w_{ij} = connection weight from input i to hidden unit j .

The logistic, or sigmoid function, is used as the activation function on hidden and output units, because its output closely resembles a threshold-step function and is differentiable; it is depicted in Equation 2

$$f(z) = \frac{1}{1 + \exp(-z)} \quad (2)$$

Next, the output calculated for each input pattern is compared to the corresponding target output, and gradient steepest descent, using the square of the difference between the target and observed output as an error function, is used to adjust the connection weights and bias thresholds so that the network output will be closer to the target output value the next time it is presented the input pattern. Because the sigmoid function only asymptotically approaches 0 or 1, (non-incident or incident), target values of 0.1 and 0.9 are used instead. Partial derivatives of the error function are taken with

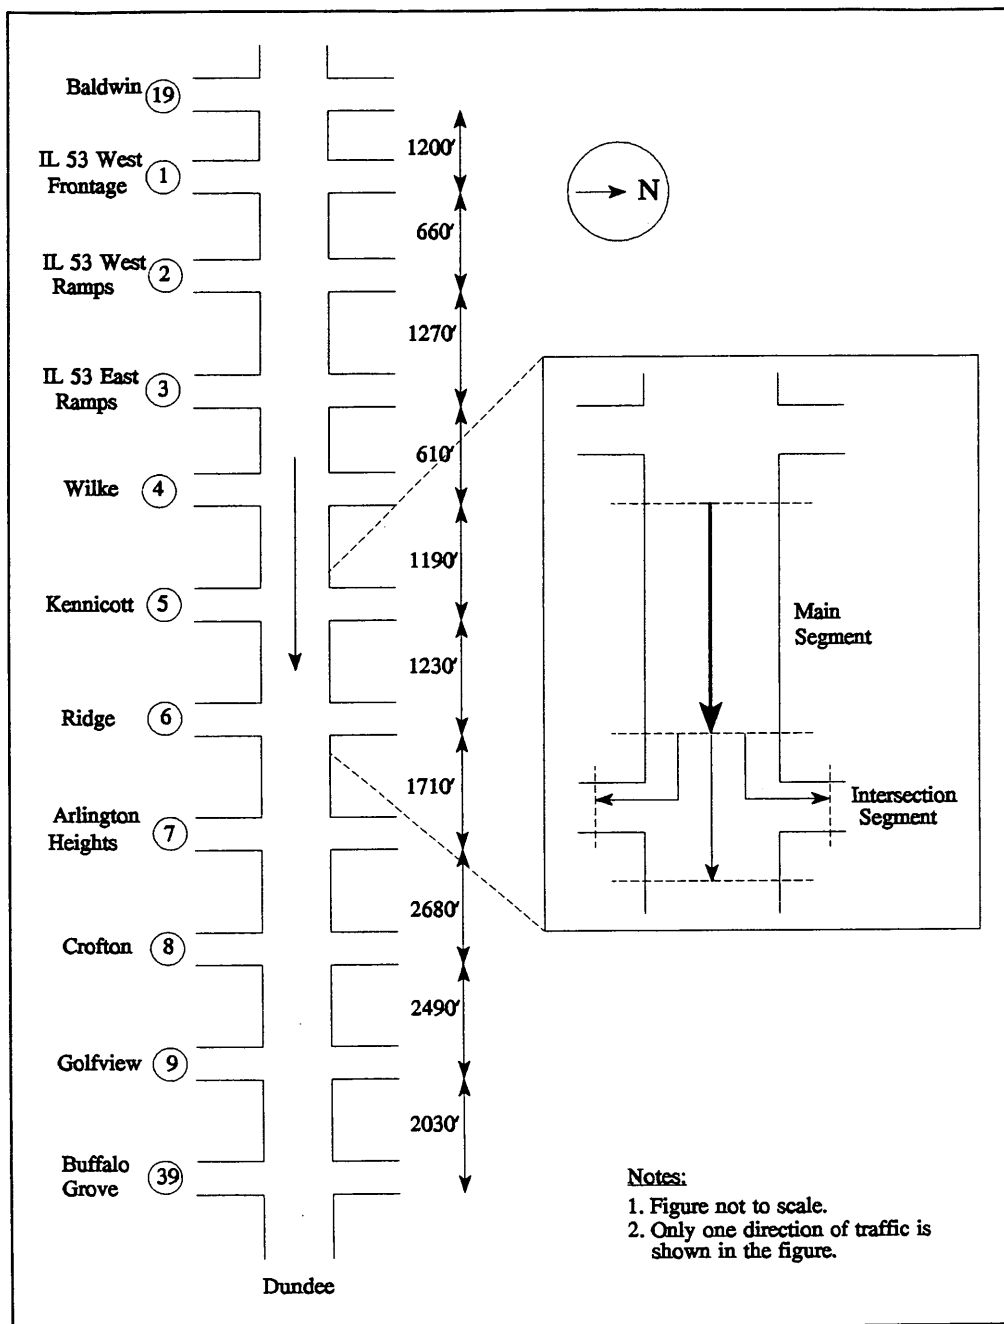


FIGURE 2 Simulation network.

respect to each connection weight and are then used to adjust the weights as expressed in Equation 3

$$\Delta w_{ij} = \eta \delta_j x_i \quad (3)$$

where

- Δw_{ij} = change computed for the connection weight from unit i to unit j ;
- η = learning rate (controls the rate at which the network makes adjustments);

δ_j = propagated backward through unit j , and
 x_i = activation of unit i .

When j is the output unit, the propagated error is given by Equation 4,

$$\delta_y = (y^* - y) f'(v_y + \sum_{j=1}^m v_j x_j) \quad (4)$$

where

δ_y = delta value for the output unit;

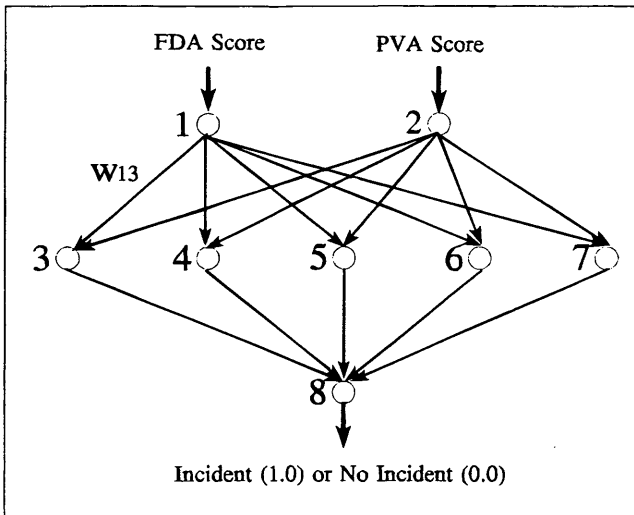


FIGURE 3 Algorithm output fusion network structure.

y^* = desired target output value;
 $f'(z)$ = first order derivative of the sigmoid function;
 x_j = activation of the j th hidden units; and remaining symbols are as previously defined.

For hidden units, Equation 5 is used

$$\delta_j = \delta_y w_{jy} f'(w_{0j} + \sum_{i=1}^n w_{ij} x_i) \quad (5)$$

where δ_j is the delta value for hidden unit j .

The learning rate determines how much to change each weight value after each input vector presentation. A better performance of the learning algorithm can be achieved by incorporating a momentum term, as shown in Equation 6

$$\Delta w_{ij}(t) = \eta(\delta_j x_i) + \alpha \Delta w_{ij}(t - 1) \quad (6)$$

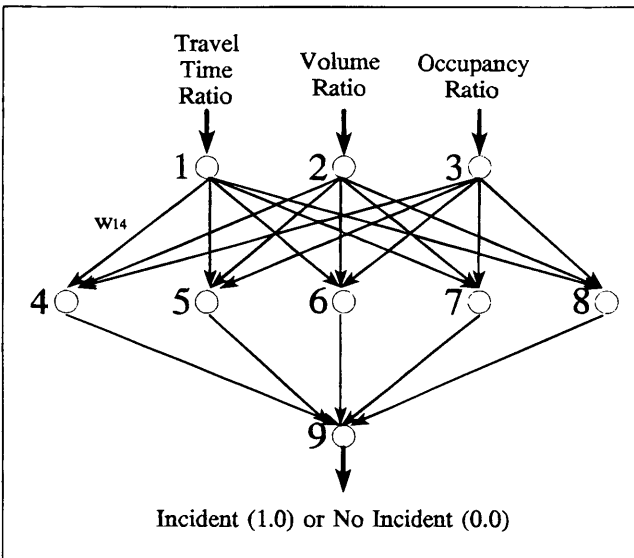


FIGURE 4 Integrated fusion network structure.

where $\Delta w_{ij}(t)$ is the weight change for the connection from unit i to unit j for presentation t , the values subscripted with t are determined for presentation t , and α is the momentum rate. The momentum term causes weight changes to move steadily in the same average direction and prevent it from settling in a local minimum. For both of these networks, a learning rate of 0.2 and a momentum rate of 0.8 were used.

The following procedure was followed for each network:

1. Ten percent of the training data files were selected randomly and set aside for testing.
2. Network training began with asymmetric connection weights randomly assigned values between -0.5 and $+0.5$ (15); the remaining data files were divided into seven groups and added to the network training set one at a time every 50 epochs.
3. Starting at 500 epochs, network performance was tested on both the training and reserved data sets at regular 250-epoch intervals; training was discontinued when root mean square error stopped decreasing on both training sets or began to increase on the reserved set (indicating overfitting of the training data) (11).

RESEARCH FINDINGS

The Algorithm Output Fusion Network was trained using the parameters and procedure just described. Figure 5 shows the relationship between the root mean square error (RMSE) of the network output (from the target values) and epoch (or iteration) as the network trained. The wide oscillations early on result from the incremental introduction of data files into the training vectors. Note that the error does not change much after 1000 epochs of training. Figure 6 shows the same plot for the Integrated Fusion Network. RMSE again drops quickly by the thousandth epoch, but instead of holding at a constant value, it continues to decrease at an extremely slow rate. However, detection performance through this period does not appreciably improve, so training was terminated.

The performance of both networks on both the reserved test data and the training data is depicted in Table 1. The following performance measures are shown:

1. RMSE over all vectors in each data set;
2. Detection rate, the proportion of known incident observations (each individual period that an incident occurred) correctly classified; and
3. The false alarm rate, the proportion of nonincident observations incorrectly classified as incidents.

The Algorithm Output Fusion Network detects all of the incident observations in the reserved data set with no false alarms, but detects only 81 percent and misclassifies 0.11 percent of the nonincident observations in the training data set. The Integrated Fusion Network did not train as well. RMSE is 0.0901 for the training data and 0.1051 for the reserved data. Although RMSE at this stage continues to decrease with the training data, detection rate on both data sets stabilized at 66 percent on the training data and 70 percent on the reserved data, so further training would not yield better results. This network resulted in an unacceptably high false alarm rate for both data sets.

It is also worth noting that the networks perform much better on the reserved test data than on the training data with which they learned. This is an unexpected result, as it is analogous to a student scoring better on questions she had not seen before than on the ones

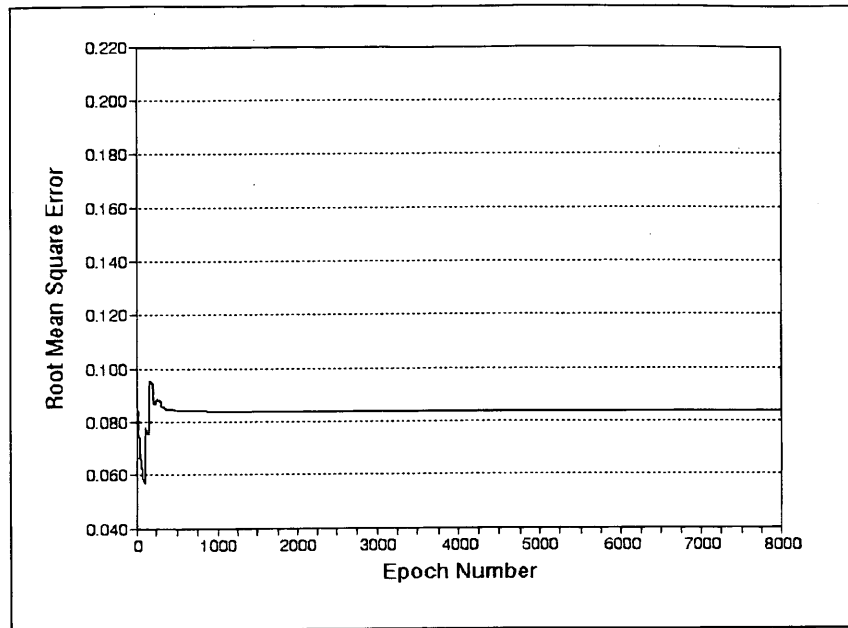


FIGURE 5 Algorithm output fusion network error by epoch.

she rehearsed prior to the examination. It turns out that the incidents in the reserved data set (though selected randomly) caused more extreme traffic conditions on average than did the incidents in the training data set. This does not implicitly invalidate this partitioning of the data files; it is simply necessary to test the network with both data sets to understand its true performance.

Table 1 also lists performance measures for the incident detection algorithms for comparison. The Integrated Fusion Network does not perform much better than the algorithms, but the Algorithm Output Fusion Network dominates all other processes with its

much greater detection rates. Data fusion process and algorithm performance can be compared more directly by plotting adjusted algorithm output scores against each other on a grid for each incident, marking each observation according to whether or not the incident was detected. These plots are shown in Figure 7 for the Algorithm Output Fusion Network and in Figure 8 for the Integrated Fusion Network.

The X and Y axes divide each plot into four quadrants. The lower left quadrant contains known incident observations which both algorithms fail to detect, and the upper right quadrant, those which




FIGURE 6 Integrated fusion network error by epoch.

TABLE 1 Neural Network Performance Summary

	RMS Error	Detection Rate	False Alarm Rate
Algorithm Output Fusion			
Training Data	0.0838	81.5%	0.11%
Reserved Data	0.0288	100.0%	0.00%
Integrated Fusion			
Training Data	0.0901	65.7%	0.54%
Reserved Data	0.1051	70.0%	0.96%
Fixed Detector Algorithm	—	65.9%*	0.00%*
Probe Vehicle Algorithm	—	53.7%*	0.00%*

— indicates value not available

* Source: (14)

both algorithms do detect. Observations in the upper left quadrant are missed by the fixed detector algorithm but detected by the probe vehicle algorithm, and those in the lower right quadrant are detected by the fixed detector algorithm but missed by the probe vehicle algorithm. This interpretation helps identify improvements offered by each data fusion process over the algorithms.

The Integrated Fusion Network offers no appreciable improvement over the incident detection algorithms, as it misses many incidents that both algorithms detected, and detects few that were missed by either algorithm. The Algorithm Output Fusion Network does much better, detecting all but two incidents that were detected by at least one of the algorithms. Even though it misses two incidents detected by the fixed detector algorithm, it detects five that the algorithm missed.

SUMMARY AND CONCLUSIONS

The results show that neural networks can be trained to detect incidents in arterial street settings at least as well as many conventional algorithms do in the less challenging freeway setting. The Algorithm Output Fusion Network detected well over 85 percent of the incidents in the data sets with no false alarms. A performance evaluation of a number of prominent conventional freeway incident detection algorithms (2) found false alarm rates over 0.5 percent associated with detection rates this high. Note, however, that a greater variety of incident types is included in these other studies and that the networks considered here were all trained with data collected from a traffic simulation rather than with field data as the conventional algorithms were. To the extent that the traffic simulation program used to generate the training data was calibrated to replicate the operation of a real street, it may well be reasonable to compare performance with the other algorithms directly. Nevertheless, confidence in this result would increase if similar results were obtained from a network tested (or trained) with field data.

It has also been shown that incident detection system performance can be improved by combining information from different data sources, in this case fixed detectors and probe vehicles. This idea is partially supported by the plots of algorithm output scores for incident records classified by whether each incident was detected or missed, but more positively by the detection rates reported on Table 1. Since the algorithms were calibrated to report no false alarms, they miss many marginal incidents which the networks are able to detect by combining the algorithm reports. This is good news for ITS demonstrations such as ADVANCE, which use information from a variety of sources. The bad news is that data from fixed detectors appear to be more reliable than data collected by probe vehicles, as evidenced by the superior performance of the fixed detector algorithm (14), and the availability of fixed detector data will be limited in ADVANCE. However, the probe vehicle incident detection algorithm was able to detect as many as 61 percent of the incidents alone with no false alarms (14), so a reason-

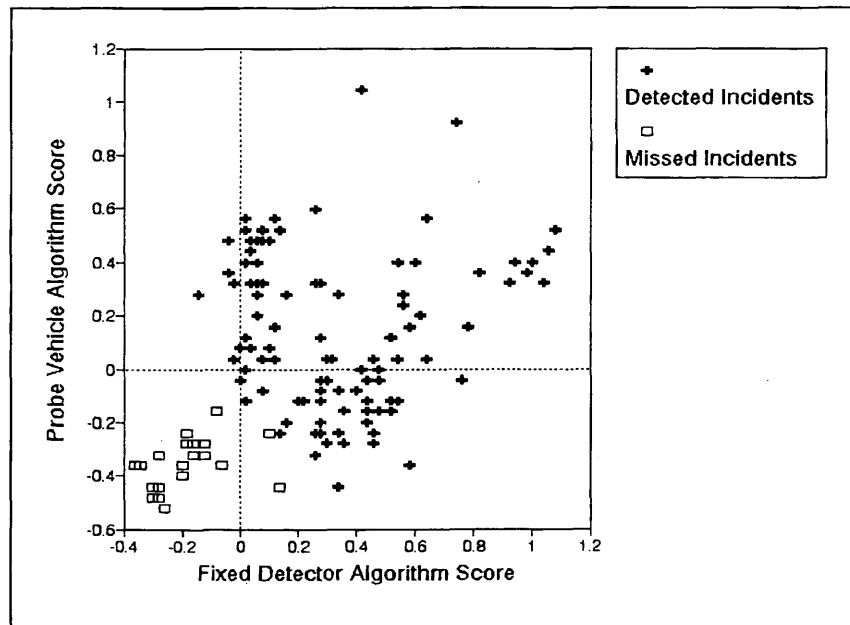


FIGURE 7 Algorithm scores for all incident observations—algorithm output fusion network.

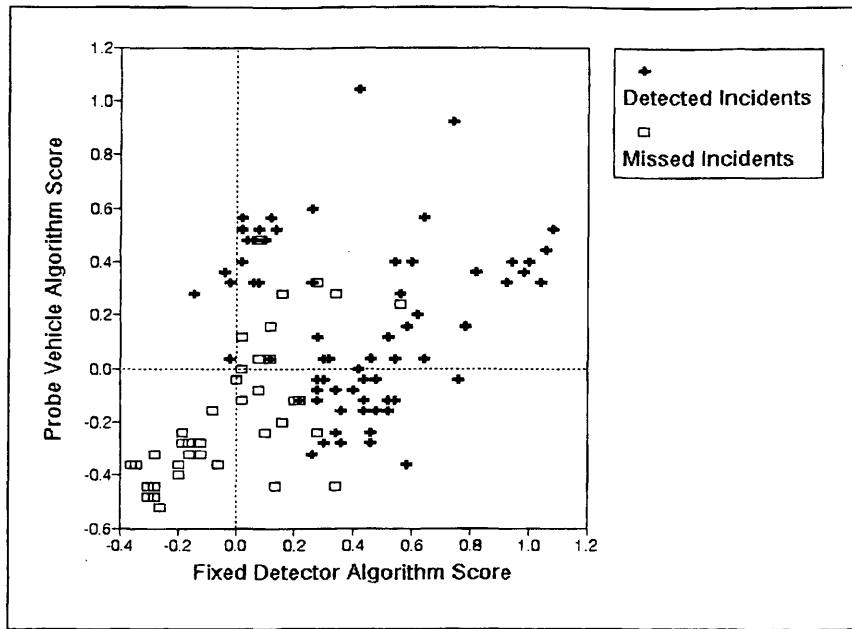


FIGURE 8 Algorithm scores for all incident observations—integrated fusion network.

able incident detection system is attainable even without fixed detector data; these data simply permit the system to detect a few more less-severe incidents.

FUTURE DIRECTIONS

A desirable objective of this (or any) incident detection system is to be an “off-the-shelf” algorithm which does not require recalibration for each implementation site. While the network performance

observed suggests that it was not necessary to learn different parameters for each highway link, this total portability feature will become much more reliable as the variety of traffic and street characteristics in the training data, and thus, the transferability of the result, increases. The ADVANCE demonstration, once it becomes operational, can provide field-collected fixed detector and probe data that implicitly include this variability. Observed travel times will be collected regularly for all links traveled by the participating probe drivers; this information could be combined with any fixed detector data available for those links.

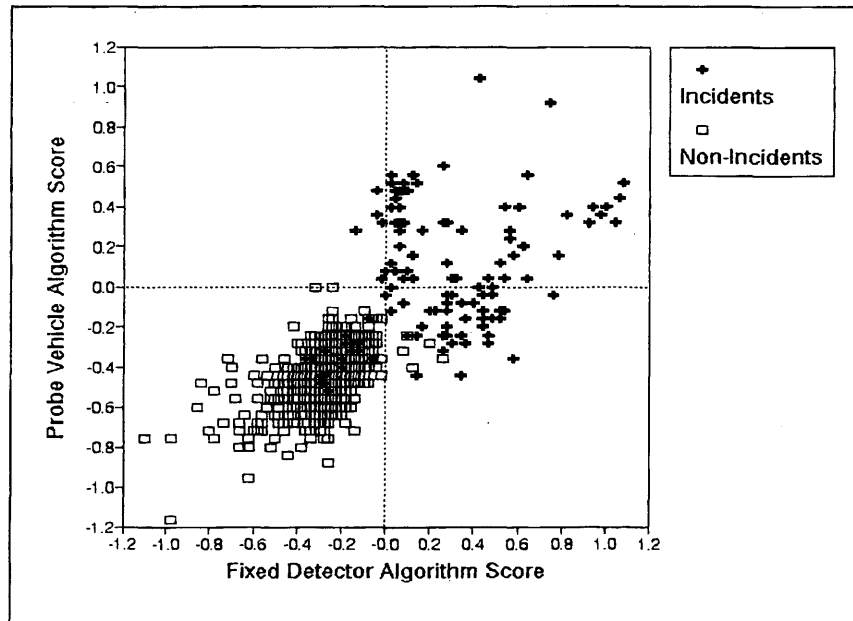


FIGURE 9 Algorithm score patterns.

A number of additional future research directions are suggested by this work. The networks should be retrained with field data, as discussed, to confirm these findings and to investigate other issues such as optimal network and input data representations and the effect of performance of using additional data sources (such as motorist calls and emergency dispatch communications).

The simulation data used for calibration of the probe vehicle incident detection algorithm and for training the neural networks considers 25 percent of the vehicle stream to be probes. The effect of much smaller probe vehicle proportions and the number of travel time reports included in each aggregation interval should be investigated.

This research is concerned with detecting incidents on arterial streets; a logical extension is development of systems that use this capability to modify traffic control parameters (e.g., signal timings) in response to observed traffic conditions. Such systems will become increasingly important as ITS implementations attempt to extract more and more capacity from existing highway networks.

ACKNOWLEDGMENTS

The information in this report was partially funded and developed for the ADVANCE Project, a joint ITS undertaking by the (FHWA) Federal Highway Administration, the Illinois Department of Transportation, the University of Illinois and Northwestern University operating under the auspices of the Illinois Universities Transportation Research Consortium, Motorola, Inc., and other participating organizations.

REFERENCES

1. Hall, R. W. Non-Recurrent Congestion: How Big is the Problem? Are Traveler Information Systems the Solution? *Transportation Research Part C*, No. 1, March 1993, pp. 89–103.
2. Stephanedes, Y. J., A. P. Chassiakos, and P. G. Michalopoulos. Comparative Performance Evaluation of Incident Detection Algorithms. *Transportation Research Record 1360*, TRB, National Research Council, 1992, pp. 50–57.
3. Han, L. D., and A. D. May. *Automatic Detection of Traffic Operational Problems on Urban Arterials*. UCB-ITS-RR-89-15, Institute of Transportation Studies, University of California at Berkeley, 1989.

4. Hounsell, N. B., M. McDonald, and C. F. S. Wong. Traffic Incidents and Route Guidance in a SCOOT Network. In *Traffic Management and Road Safety, Proceedings of Seminar B, Planning and Transport Research and Computation (International) Co. Ltd. (PTRC) Transport and Planning Summer Annual Meeting*. University of Bath, England, September 12–16, 1988.
5. Thancanamootoo, S., and M. G. H. Bell. *Automatic Detection of Traffic Incidents on a Signal-Controlled Road Network*. Transport Operations Research Group, University of Newcastle upon Tyne, Research Report No. 76, 1988.
6. Boyce, D. E., A. Kirson, and J. L. Schofer. Design and Implementation of ADVANCE: The Illinois Dynamic Navigation and Route Guidance Demonstration Program. In *Vehicle Navigation and Information Systems Conference Proceedings*, Society of Automotive Engineers, Dearborn, Michigan, October 20–23, 1991, P-253 Part 1.
7. Khattak, A. J., J. L. Schofer, and F. S. Koppelman. Factors Influencing Commuters' En Route Diversion Behavior in Response to Delay. In *Transportation Research Record 1318*, TRB, National Research Council, 1991, pp. 125–136.
8. Prosser, N. A., and S. G. Ritchie. Real-Time Knowledge-Based Integration of Freeway Surveillance Data. Presented at the Transportation Research Board 70th Annual Meeting, January 13–17, 1991.
9. Sumner, R. Data Fusion in Pathfinder and TravTek. In *Vehicle Navigation and Information Systems Conference Proceedings, P-253 Part 1*, Society of Automotive Engineers, Dearborn, Michigan, October 20–23, 1991.
10. Hamburger, H. Representing, Combining and Using Uncertain Estimates. *Uncertainty in Artificial Intelligence*, (L. N. Kanal and J. F. Lemmer, eds.) Amsterdam: Elsevier Science Publishers B. V., 1986.
11. Caudill, M. Neural Networks Primer, Parts I, II and III. Reprinted from *AI Expert*, Miller Freeman Publications, San Francisco, 1990.
12. Rumelhart, D., and J. McClelland. *Parallel Distributed Processing Volume I: Foundations*. MIT Press, Cambridge, Mass., 1986.
13. Goldblatt, R. B. *Development and Testing of INTRAS, A Microscopic Freeway Simulation Model, Vol. 3, Validation and Application*. FHWA Offices of Research and Development, Traffic Systems Division, 1980.
14. Koppelman, F. S., J. L. Schofer, N. Bhandari, V. Sethi, and J. N. Ivan. *Calibration of Probe and Fixed Detector Algorithm Parameters with Simulated Data*. Evanston, Illinois: ADVANCE Project Technical Report TRF-ID-151, The Transportation Center, Northwestern University, 1994.
15. Hertz, J., A. Krogh, and R. G. Palmer. *Introduction to the Theory of Neural Computation*. Addison Wesley, 1991.

The contents of this report reflect the views of the authors who are solely responsible for the information and accuracy of the data presented. It does not constitute a standard, specification, or regulation. The contents of this report do not necessarily reflect the official policy of any of the ADVANCE parties, nor does any mention of manufacturers or products constitute an endorsement by the ADVANCE parties.

Publication of this paper sponsored by Committee on Artificial Intelligence.