Final report of ITS Center project: FleetForward traffic forecasting]

UVA Center for Transportation Studies

A Research Project Report

For the Center for ITS Implementation Research

A U.S. DOT University Transportation Center

AN INVESTIGATION INTO INCIDENT DURATION FORECASTING FOR FLEETFORWARD

Principal Investigator: Brian L. Smith

Graduate Assistant: Kevin W. Smith

Center for Transportation Studies
University of Virginia
Thornton Hall
351 McCormick Road, P.O. Box 400742
Charlottesville, VA 22904-4742
804.924.6362

August, 2000
UVA-CE-ITS_01-5

Disclaimer

TABLE OF CONTENTS

INTRODUCTION

Traffic condition forecasting is the process of estimating future traffic conditions based on current and archived data. Real-time forecasting is becoming an important tool in Intelligent Transportation Systems (ITS). This type of forecasting allows ITS to enact control and management strategies that are "one step ahead" rather than "one step behind" the onset of traffic conditions (Williams and Smith, 1999). For example, an ITS traffic management system can take measures to anticipate congestion rather than reacting to congestion once it is present. Real-time forecasting has benefits to many research fields including route guidance, incident management, public transportation operation, and traveler information (Perrin and Martin, 1998).

The most common traffic conditions that are forecasted on a real-time basis are flow rate and travel time. The specific traffic condition that the University of Virginia's Smart Travel Laboratory is attempting to forecast in this research effort is incident duration, a relatively new area of research for transportation forecasting. To date, there has been limited research into models that can predict how long a certain incident will affect traffic.

It has been said that the target audiences of predictive traffic information are commuters and motorists on business (Al-Deek, et al., 1999). Motor carriers fit nicely in this category, as their business is to provide transportation services. Incident duration forecasts will be extremely important to motor carriers and thus will be a useful tool for FleetForward, a traveler information system for motor carrier operations. Knowing how long an incident will affect traffic allows motor carrier dispatchers and drivers to more intelligently schedule and route shipments.

FLEETFORWARD

FleetForward is an operational test designed to demonstrate the impact of real-time traffic information on commercial vehicle operations, such as dispatching and routing. The test was initiated in 1997 by the American Trucking Association (ATA) Foundation as a public-private partnership involving 14 government agencies, private technology firms, and representatives of the motor carrier industry. FleetForward incorporates real-time traffic data from SmartRoute's SmarTraveler system and the I-95 Corridor Coalition's Information Exchange Network (IEN) with a traditional, "static" routing and scheduling tool.

The difference between these two traffic data sources lies in their respective scopes. SmartRoutes provides relatively high-resolution metropolitan traffic data for a number of cities along Interstate 95, including Washington, Boston, Philadelphia, and New York. The data includes the location of highway incidents such as accidents and work zones, along with link travel times for many arterial roads. On the other hand, the IEN data has a larger scope. The I-95 Corridor Coalition consists of state and transportation agencies along the I-95 corridor from Virginia to Maine. The IEN serves as a mechanism for states in the Coalition to share information about major, corridor-level, incidents. Therefore, archival data from the IEN contains major highway incidents along the entire length of I-95.

IEN Data Quality

The IEN incident database contains a large amount of data major incidents that occurred in the I-95 corridor from 1997 - 1999. When a traffic incident is reported to the IEN, the agency sends a report marked NEW. When the incident has ended the same agency sends a report marked CLOSE. All of the incident characteristics are included in the NEW report, but not the CLOSE report. In terms of analyzing past incidents, the sole function of the CLOSE report is to calculate the duration of the incident. Thus, an incident in the database can not be used for analyzing duration without the presence of a NEW and CLOSE report. Table 1 above shows that of the 8166 incidents in the IEN database, only 7235 incidents are available for analysis due to the lack of both a NEW and CLOSE report.

Table 1. Summary of incident reports from IEN database.

Year	Total Incidents	Incidents Without a NEW report	Incidents Without a CLOSE report
1997	5441	109	517
1998	2623	66	142
1999	102	0	97
1997 – 1999	8166	175	756

There are some interesting trends in the IEN database in terms of the time and location of the highway incidents. The number of incidents reported for each month is shown in Figure 1.

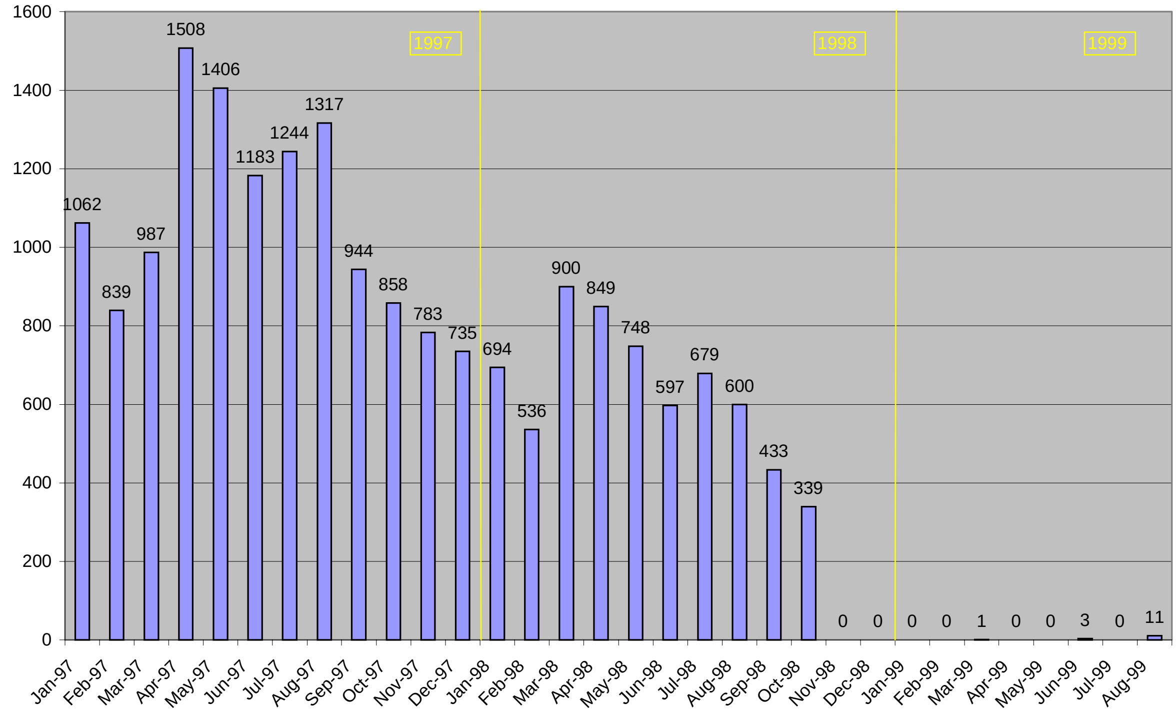

Figure 1. The number of reported incidents by month.

This shows an interesting trend in the occurrence of incidents in the IEN database. The cause for the general decline in the number of incidents is unknown and raises questions about the completeness of the IEN database. A second trend is the location of the traffic incidents. Figure 2 shows the percentage of total incidents that occurred in each state.

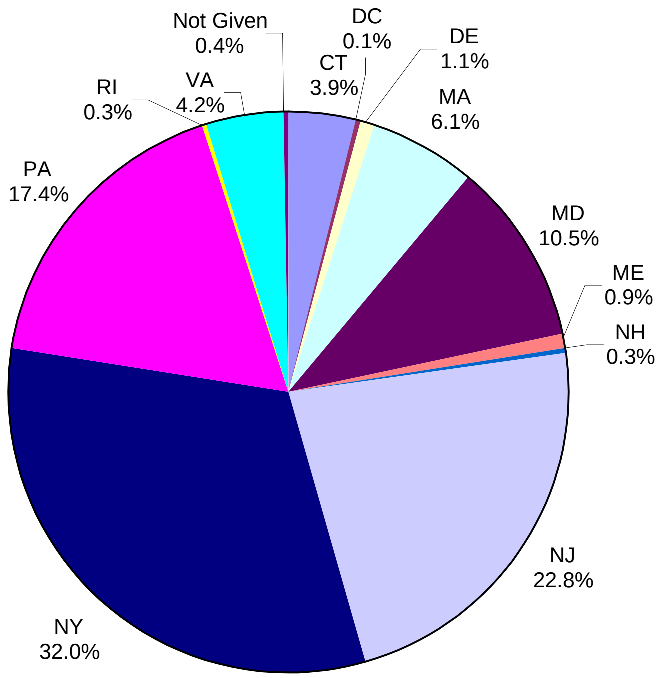

Figure 2. The breakdown of all incidents from 1997-1999 by state.

Figure 2 shows that over half of the reported incidents occurred in the New York and New Jersey section of the I-95 Corridor. It is unclear if this dominance is due to more actual incidents or more reported incidents due to other factors such as the number of NY/NJ agencies, duplicate incident reports, or reporting of minor incidents that would not be reported by other agencies.

One piece of data that we expected to benefit our analysis of incident duration was the expected duration data that was recorded for each incident. However the entries in this field of the database are not consistent and discernable. Most of the values in this field are in the format “12/30/1899 00:15:00” which may indicate an expected duration of 15 minutes. However, values also took other formats that are not as easy to interpret. For example, the database contains estimated duration entries such as “01/01/1900 00:00:00,” “04/09/1900 00:00:00,” “8.33333333 E –02,” and 0.0625. It appears that there was no standard procedure for entering the estimated duration in an IEN report and thus the data does not provide any useful information for our incident analysis.

Overall, there are some questions regarding the quality of the IEN incident data. We have no reason to doubt the data on incident characteristics such as the incident type, location, time, and lane closure. The main question is the accuracy of the incident duration since this calculation requires two separate reports to be sent from the same agency. Considering that a number of incidents were never closed leads us to believe that other incidents might not have been closed at the appropriate time in the IEN database.

Delivery Techniques

An important aspect of FleetForward is the two different delivery approaches of real-time traffic data to motor carriers. One delivery technique is using the Internet to reach the motor carriers. FleetForward has developed a webpage that displays real-time traffic data for a number of metropolitan areas and the entire I-95 corridor. The webpage displays a map with color-coded highway segments based on their level of congestion (see Figure 3).

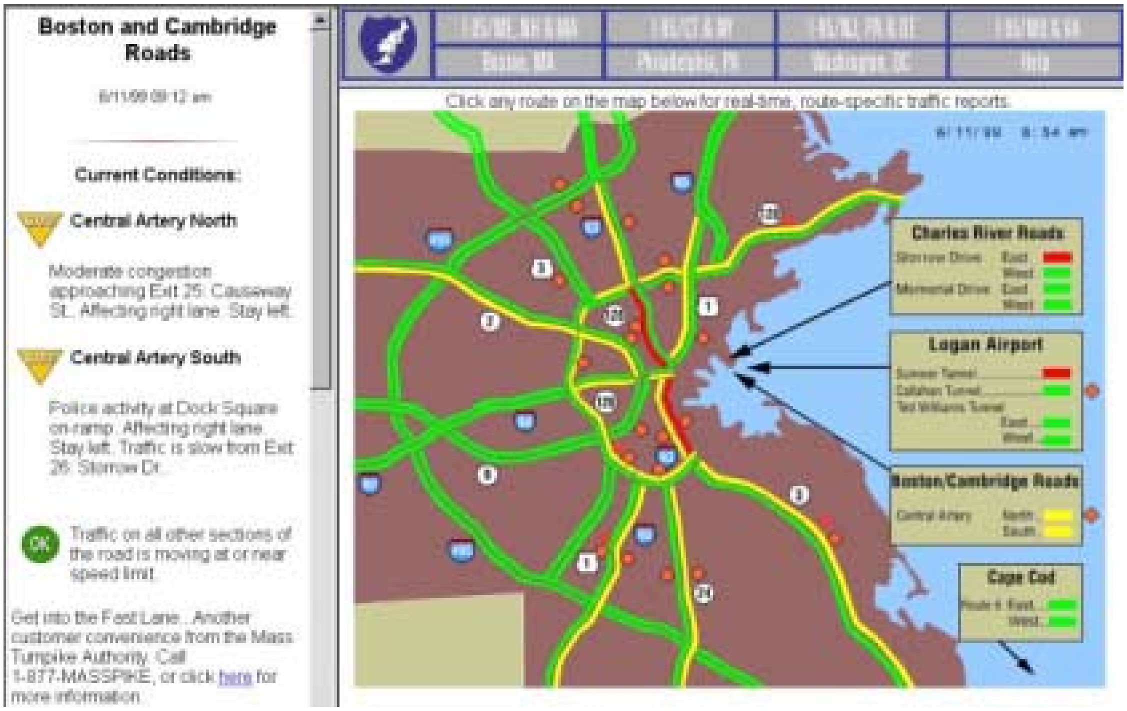

Figure 3. Screenshot from the FleetForward webpage.

This approach is graphical and allows for easy interpretation of areas of congestion. Many motor carrier dispatchers also use existing software packages to efficiently manage their fleet. These programs have large street databases that are used to calculate the most appropriate route from origin to destination. FleetForward has extended an existing software package in order to incorporate real-time traffic data into routing decisions.

A key limitation of the FleetForward system that is incorporated into the routing/scheduling software is how highway incidents are handled by the routing procedure. If an incident has occurred on a particular road, the incident is marked on the map and that particular link is disqualified from routing consideration. The negative impact of this approach is that the link with the incident may be unnecessarily removed from consideration. For example, consider the case when the software is used to find the best route from Boston to New York City. A dispatcher runs the algorithm while a portion of I-95 has reported an accident. The software would then reroute the motor carrier off of the interstate and around the incident. However, this does not take into account the fact that the driver may not reach that segment for another several hours, and the accident may be cleared by that point with traffic flow returning to normal. What is needed in this case is the expected duration of the highway incident to facilitate effective routing and dispatching decisions.

In summary, FleetForward has proven successful as the first operational test to merge multiple traffic data sources and deliver an information stream via web-based traffic maps and integrated routing and dispatching software to motor carriers for the purpose of improving fleet management decision support. Participating motor carriers identified numerous benefits to their operations from FleetForward including improved on-time performance, greater customer and driver satisfaction. They also indicated several enhancements that would increase the value of FleetForward as a management tool, including the ability to predict incident duration. The remainder of this report focuses on the research effort to develop an incident duration forecasting capability.

INCIDENT DURATION FORECASTING

As seen above, incident duration forecasting is needed in order to improve the usefulness of advanced commercial vehicle operations tools, such as FleetForward. In our research effort, we attempted to use a large archived database of past highway incidents to find patterns and relations that would allow for the forecasting of current incident duration. The IEN database used by FleetForward was utilized in the research effort. The following characteristics for each incident were used in evaluating forecasting models.

- A unique ID Number for each incident
- The type of highway incident (accident, construction, debris, etc.)
- The time and date of the incident
- The incident duration
- The location of the incident (state)
- The number of lanes closed during the incident

These characteristics can be used as independent variables to define the state of an incident for forecasting duration. There are many methods and models that can be used to forecast duration.

Past techniques used to predict incident duration have ranged from statistical to heuristic approaches. Standard regression models have the advantage of being easily understood, but tend to oversimplify the representation of an incident (Nam and Mannering, 2000). Probabilistic approaches such as lognormal distributions and analysis of variance have been used with success to analyze incident duration (Nam and Mannering, 2000). A new approach to statistically evaluate the factors that tend to influence incident duration is hazard-based models, a technique that has been used in the past to analyze traveler activity behavior, automobile ownership, and traffic queuing (Nam and Mannering, 2000).

Nonparametric Regression

The forecasting approach explored in this research for use in FleetForward is nonparametric regression. Nonparametric regression is a forecasting technique that requires no strict assumptions regarding a functional relationship between dependent and independent variables. Unlike traditional regression models that define a relationship for all ranges of dependent variables, nonparametric regression focuses on a specific area, or neighborhood, of past system states that are similar to the current system state. The past instances in this neighborhood are then combined (usually a weighted average) to predict the dependent variable value. This method relies heavily on a having a wide range of quality data to make predictions (Smith, et al., 2000).

The key to an effective nonparametric model is effectively defining a neighborhood of past instances. The two most popular approaches are kernel and nearest neighbor (Altman, 1992). A kernel neighborhood is defined as having a constant bandwidth on the independent variable space (Smith, et al., 2000), centered on the current state under investigation. A nearest neighbor neighborhood is defined as having a constant number of data points that includes those “nearest” to the current system state. The main difference between these two approaches is that the nearest neighbor guarantees that a prediction is made, while a kernel neighborhood may not be able to find any past similar instances within the predefined bandwidth.

As the name implies, in order to define “near” neighbors, an appropriate distance metric in the state space must be defined. Often an appropriate choice is Euclidean distance. This is most applicable to systems with numerical inputs and outputs. Other distances metrics can use weighted distances in the system state space (Smith, et al., 2000). The choice of the distance function depends on the nature of the data and the experiences of the developers.

Once a neighborhood has been defined, a prediction is generated. The most common prediction generation approach is to use the average of the dependent variables for the selected neighbors. Another popular method is weighted average, where nearer neighbors are given a larger weight in the prediction. This area of nonparametric regression is rapidly expanding with an array of new methods being tested (Smith, et al., 2000).

RESULTS

For this project, a simple nonparametric regression algorithm was developed that used an unweighted average of a kernel neighborhood. The independent variables used in the experiment are listed in Table 2. The IEN incident database was randomly split into incidents for model development and testing. For a wide range of kernel sizes, there were 1085 incidents from 1997 to

1999 tested. The two main measures of effectiveness were the mean absolute percent error (MAPE) and the number of predictions made by the model. Since this experiment used a kernel neighborhood, it was possible that a small kernel would result in the model being unable to find any past incidents within that neighborhood size. The MAPE is simply the mean of the percent errors for the 1085 test incidents for a given kernel. The percent error in this case is defined as the ratio of difference between the predicted and actual incident duration and the actual duration.

Table 2. Independent variables used in nonparametric regression model.

Independent Variable	Possible Values	
Type of Incident	Accident	Debris in Road
	Disabled Vehicle	Hazardous Material (HAZMAT)
	Truck Incident	
Time of Day	AM Peak	Mid-day
	PM Peak	Off-hour
Day of Week	Weekday	Weekend
Location	Virginia	D.C.
	Maryland	Delaware
	Pennsylvania	New Jersey
	New York	Connecticut
	Rhode Island	Massachusetts
	New Hampshire	Maine
Percent of Lanes Closed	No lanes	< 50% of Lanes
	> 50% of Lanes	All Lanes

An examination of the table above reveals that there are 27 possible incident characteristics. The single dependent variable for use in the model is the duration of the incident. We decided to combine 27 characteristics into a single independent variable to be used in the nonparametric regression. To do this each of the 27 characteristics was represented as a binary code to indicate their presence or absence from a particular incident. Then a penalty constant was assigned for each variable. Thus, each incident was given a penalty function that was the product of the binary matrix and the penalty constant matrix, as follows:

$$Y = \sum_i X_i \, C_i$$

where Y = total penalty of incident

Xi = 1 if variable is present, 0 if not

Ci = penalty for the ith independent variable

The penalty constants that were used are presented in Table 3.

TABLE 3. Penalty values of forecasting variables.

		Independent Variable (binary)	Penalty Variable	Value
Type of Incident	Accident	X_1	C_1	20
	Debris in Road	X_2	C_2	40
	Disabled Vehicle	X_3	C_3	60
	HAZMAT	X_4	C_4	80
	Truck Incident	X_5	C_5	100
Time of Day	AM Peak	X_6	C_6	0.01
	Mid-day	X_7	C_7	0.02
	PM Peak	X_8	C_8	0.03
	Off-hour	X_9	C_9	0.04
Day of Week	Weekday	X_{10}	C_{10}	0.05
	Weekend	X_{11}	C_{11}	0.06
State	Virginia	X_{12}	C_{12}	1
	D.C.	X_{13}	C_{13}	2
	Maryland	X_{14}	C_{14}	3
	Delaware	X_{15}	C_{15}	4
	Pennsylvania	X_{16}	C_{16}	5
	New Jersey	X_{17}	C_{17}	6
	New York	X_{18}	C_{18}	7
	Connecticut	X_{19}	C_{19}	8
	Rhode Island	X_{20}	C_{20}	9
	Massachusetts	X_{21}	C_{21}	10
	New Hampshire	X_{22}	C_{22}	11
	Maine	X_{23}	C_{23}	12
Percent of Lanes Closed	None	X_{24}	C_{24}	0.1
	< Half	X_{25}	C_{25}	0.2
	> Half	X_{26}	C_{26}	0.3
	All	X_{27}	C_{27}	0.4

The purpose of the penalty constants is to assist in defining the neighborhood of an incident. The large values given to the incident type variable constrain the neighborhood search to only include one type of incident. This is logical, as it is misleading to compare a disabled vehicle incident with a major highway accident. The state variables also have large values for the reason that state agencies differ in their approach to clearing incidents. Also, the state variables are arranged in the order that I-95 travels through the Northeast. So a neighborhood that includes other states will include neighboring states first.

The choice of an appropriate kernel size was guided by an empirical approach. For this experiment we tested a wide range of kernel sizes, from very small kernels that forced exact matches of the independent variables to large kernels where all past incidents were considered in the neighborhood. Figure 4 shows a range of kernel sizes where the best results were found. This chart shows the MAPE (between 100% and 120% error) and the number of predictions returned for a given kernel. A general trend is that the percent error of the prediction decreases rapidly to a lower limit and then steadily increases as the kernel size increases.

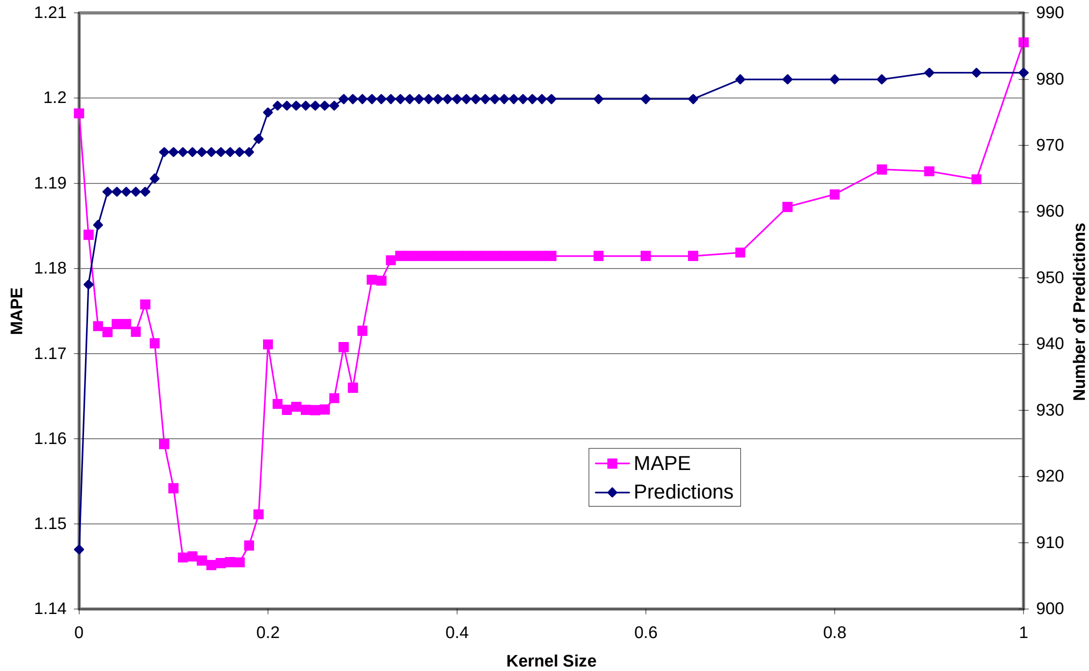

Figure 4. Results from nonparametric regression model.

CONCLUSIONS

The results illustrated in Figure 4 show that the predictions of incident duration from this model differ from the actual incident duration by an average of over 100%. This error is unacceptably large for a forecasting model to be used in the field. While the research team has identified a number of areas to improve the implementation of the nonparametric regression approach, this is not likely the driving factor. It is likely that this error is primarily attributable to the database used, and specifically the choice of independent forecasting variables.

The IEN database contains a significant amount of data describing each traffic incident. This data is very important to support communication among transportation agencies, but may not be suited for an archived database being used for incident forecasting. The independent variables used in this experiment were time of day, day of week, incident type, location, and lane closure. The time of day and day of week are useful variables to define where the incident occurred in the normal cycle of daily traffic. The location variable with the state where the incident occurred was chosen to represent the types of assistance given to an incident. It is possible that all states along the I-95 Corridor have similar response plans to highway incidents and the same personnel and procedures are used. A more representative variable for forecasting would be the specific assistance given to the incident, such as towing, pushing vehicle off road, fire department or police response, medical attention, or other types of assistance. The lane closure variable was chosen to show the severity of the incident. More preferable variables include the number of vehicles involved, personal injuries, presence of trucks, damage to roadway, and other incident severity characteristics.

From a statistical standpoint, the independent variables used in this experiment may not be significant to incident duration. Table 4 provides statistics for the IEN database in terms of a single characteristic. This table shows that when all of the incidents are broken down by time of day, each category has a similar average duration and standard deviation. The variable that shows the largest range of duration for each possible value is the incident type, an expected independent variable for any incident forecasting model. It can be argued that the incident type should not be used as an independent variable, but that the different incident types should be clustered. For example, this would avoid defining the neighborhood of an accident with a past instance of a disabled vehicle incident.

Table 4. Statistical summary of some common incident characteristics.

Variable	Value	Number of Data Points	Sample Duration Mean	Sample Duration Standard Deviation
Time of Day	AM Peak	605	75.9	53.7
	Midday	902	74.8	53.8
	PM Peak	818	70	51.7
	Off-hour	473	75.6	52.8
Day of Week	Weekday	2474	73.1	52.4
	Weekend	324	78.9	57.1
Type of Incident	Accident	156	74	53.6
	Accident, Hazmat	26	88.3	53.7
	Accident, Lane Closed	1033	60.6	43.2
	Accident, Multi Vehicle	33	66.8	44.2
	Accident, Road Closed	351	85.4	54.2
	Debris In Road	46	69.6	49
	Disabled Vehicle	112	44.2	35.4
	Jack Knifed TT	20	87.9	54.5
	Misplaced Truck	36	92	53.1
	Overturned TT	78	120	53.1
	Overturned Vehicle	64	71.2	48.9
Percent of Lanes Closed	0	1350	70.5	53.6
	25	55	74.7	54.6
	33	205	67.5	50.9
	50	224	66.2	46.1
	66	129	73.1	53.8
	75	4	84.1	21.2
	100	825	82.7	53.4
Location of Incident	Connecticut	130	70.5	49.7
	District of Columbia	2	52.6	5.74
	Delaware	39	88	57.7
	Maine	25	50.1	46.4
	Maryland	177	73.7	56.7
	Massachusetts	193	78.8	51.1
	New Hampshire	6	35.2	30.3
	New Jersey	638	86.2	57
	New York	916	66	48.8
	Pennsylvania	567	69.5	51.8
	Rhode Island	11	82.4	68.1
	Virginia	95	88.8	53.7

The statistical summary of the independent variables also shows a large standard deviation in duration for most incident characteristics. Note that in many cases, the standard deviation is nearly equal to the mean. The scattered nature of the data points is reflected in the poor percent errors for this forecasting model.

This research effort illustrates that incident duration models are of great importance to improving advanced motor carrier information systems, such as FleetForward. However, it also demonstrates that the development of an accurate incident forecasting model is quite challenging. In particular, there is a need to collect data with more descriptive incident characteristics to be used for future duration forecasting development efforts

REFERENCES

Al-Deek, H.M., D'Angelo, M.P., and Wang, M.C. (1999). Travel Time Prediction for Freeway Corridors. Transportation Research Board, 78th Annual Meeting.

Altman, N.S., 1992. An Introduction to Kernel and Nearest-Neighbor Nonparametric Regression. The American Statistician, Vol. 46, No. 3, pp. 175-185.

Nam, D., and Mannering, F. (2000). An Exploratory Hazard-based Analysis of Highway Incident Duration. Transportation Research Part A 34 (2000), pp. 85-102.

Perrin, H.J., and Martin, P.T. (1998). On-line Comprehensive Traffic Flow Estimation From Link Flow Detectors: Model Examination. Transportation Research Board, 77th Annual Meeting.

Smith, B.L., and Williams, B.M. (1999). ITS Data: Tapping the Resources for System Operation. ITS '99 Conference Paper.

Smith, B.L., Williams, B.M., and Oswald, R.K., 2000. Comparison of Parametric and Nonparametric Models for Traffic Flow Forecasting.