Ministry of Transportation
Provincial Highways Management Division Report
Highway Infrastructure Innovation Funding Program

Modeling Winter Maintenance Activities Using Classification Trees

Final Report

June 2008

HIIFP-074

Technical Report Documentation Page

Publication Title	**Modeling Winter Maintenance Activities using Classification Trees**

Author(s) and originating offices

Jean Andrey
Department of Geography
University of Waterloo, ON, Canada, N2L 3G1
519-888-4567 x33629
jandrey@fesmail.uwaterloo.ca
Corresponding Author

Alexander Brenning
Department of Geography
University of Waterloo, ON, Canada, N2L 3G1
brenning@fesmail.uwaterloo.ca

Brian Mills
Adaptation and Impacts Research Division
Environment Canada
c/o University of Waterloo, Waterloo, ON, Canada, N2L 3G1
Brian.Mills@ec.gc.ca

Denis Kirchhoff
Department of Geography
University of Waterloo, ON, Canada, N2L 3G1
dkirchhof@fesmail.uwaterloo.ca

Ammar Abulibdeh
Department of Geography
University of Waterloo, ON, Canada, N2L 3G1
aabulibd@fesmail.uwaterloo.ca

Sponsors

Ministry of Transportation of Ontario (MTO) / Highway Innovation Infrastructure Funding Program (HIIFP)

Publication Date

June, 2008

Ministry Contact

Max Perchanok
Maintenance Research Coordinator
Highway Standards Branch, Ontario Ministry of Transportation
301 St. Paul Street, St. Catharines, Ontario L2R 7R4
Phone: (416) 235-4680, Fax: (416) 235-4020
Max.Perchanok@ontario.ca

Abstract

In northern latitudes, road salt remains a key element of winter maintenance operations. At the same time, there is increased pressure to reduce salt usage without compromising level of service or safety. Road salt usage models provide a way of benchmarking and understanding spatial-temporal variations in maintenance operations, and therefore have value in working toward improved salt management practices. The current study outlines a new approach for modeling road salt usage that addresses many of the limitations of past models. This approach is developed and illustrated using automatic vehicle locator data for three seasons and one provincial highway patrol near Ottawa, Canada. Using categorical, hourly salt application rates as the dependent variable, and various sources and types of forecast and observed weather conditions as the independent variables, five different treatment modes are modeled using classification trees and various types of "indirect" modeling. Results are promising in terms of both the accuracy of predictions and the ability of this inductive approach to identify key explanatory variables and related threshold values that affect the probability of different treatment options.

Key Words

Winter index, road salt, AVL, classification trees

Distribution

Ontario
1V

Ministry of Transportation
Engineering Standards Branch Report

HSB-039

Modeling Winter Maintenance Activities using Classification Trees

June 2008

Prepared by
Engineering Standard Branch Executive Office
Ontario Ministry of Transportation

301 St. Paul Street,
St. Catharines, Ontario, Canada L2R 7R3
Tel: (905) 704-2638; Fax (905) 704-2626

Table of Contents

Introduction

WHY DEVELOP WINTER SEVERITY MODELS FOR ONTARIO HIGHWAYS?

Winter weather creates mobility challenges for virtually all northern jurisdictions. Each year road maintenance authorities in Canada spend in the order of 1.3 billion dollars on snow removal and other winter maintenance activities (Jones, 2003; Morin and Perchanok, 2003). As well, the science of snow control is evolving quickly, and increasingly involves private firms who enter into contractual arrangements with governments.

Various agencies produce documents outlining best practices in winter road maintenance and salt management. However, in order to further improve winter maintenance, there is increasing interest in the development of winter indices that rate the severity of weather conditions in a given time period with respect to resource use (salt, sand, equipment or personnel hours).

Atmospheric/weather indices are widely used. Within the road transport sector, freezing and thawing indices are considered in highway design, in modeling pavement deterioration and in setting spring load restrictions on secondary roads. In each case, the challenge is to identify how specific atmospheric conditions translate into a particular risk and/or prompt a particular response.

For the past 25 years, there has been interest in identifying different weather conditions that trigger specific maintenance responses with a long-term view of developing improved winter weather indices that are specific to road maintenance activities. This is challenging work because of the many factors that affect maintenance decisions. These extend beyond weather to include terrain, road and traffic characteristics. As well, road maintenance activities are essentially behavioural responses to predicted weather—and thus one would expect there to be sometimes different, but possibly equally appropriate, responses to the same situation.

EXISTING INDICES AND THEIR LIMITATIONS

Several winter severity indices have already been developed. Most of the work has been conducted by or for road authorities in North America and northern Europe. Examples of related indices and publications are listed below:

- PennDOT winter severity index (Rissel and Scott, 1985)
- Cost 309 index (Knudsen, 1994; Gustavsson, 1996)
- GAB index (Gustavsson, 1996)
- MOORI index (Johns, 1996)
- Strategic Highway Research Program (SHRP) index (Boselly et al., 1996; Andrey et al., 2001; Decker et al., 2001)
- Hulme Index (Cornford and Thornes, 1996; Andrey et al., 2001)
- Finnish Meteorological Institute Index (Venalainen, 2001)
- Salt day index (Andrey et al., 2001)
- Wisconsin winter severity index (Adams, 2001)
- NORIKS index (Mahle et al., 2002)
- Indiana DOT winter severity index (McCullouch et al., 2004)
- Iowa DOT winter severity index (Carmichael et al., 2004)
- Nixon and Qui's (2005) storm typology
- TAC's winter severity index models for Canadian roads (Suggett et al., 2006 and 2007)

While the list of winter indices is extensive, the progress to date has actually been quite limited. Indeed, most current models suffer from four main deficiencies:

1. Model fit is typically low for small spatial and temporal units of analysis, and is best when modeling seasonal variations in expenditures or salt use for relatively large

areas. This limits the usefulness of these indices and suggests only a limited understanding of winter maintenance decision making.

2. Model parameters may be biased because some of the assumptions of the statistical techniques are violated, e.g., independence of observations across jurisdictions, independence of weather variables, and lack of normally distributed residuals The multi-collinearity between weather variables is particularly problematic. For example, because of the importance of snowfall amount/duration in predicting salt use, temperature and wind variables, which are somewhat related to snowfall conditions, appear unimportant in most linear models.

3. Most researchers employ some form of linear modeling where the explanatory terms are additive. Often, however, responses are non-linear, data transformations are not well justified and explanatory variables are interactive or conditional in their influence. For example, precipitation amount only plays a role below some critical temperature.

4. Most models are calibrated with weather observations from national weather services. This creates two problems. First, since maintenance decisions are based primarily on predicted weather and given the probabilistic nature of weather forecasting, it is reasonable to assume that some of the unexplained variance in winter maintenance activities is associated with the mismatch between predicted and observed weather. Second, weather conditions vary locally, and site measurements—particularly those at some distance from the roads under study—may not be representative of local road conditions. Road weather information systems are beginning to remedy this situation for some variables, but the most important variables, which are related to winter precipitation type and amount, are not well captured by these sensors.

There is thus a need to explore the potential benefit of alternate modeling approaches, as well as the usefulness of both forecast and observed weather data from multiple sources

Ongoing research by MTO on winter indices

In winter, 2005, the Ministry of Transportation of Ontario commissioned Synectics Transportation Consultants to lead a study entitled, "Operational index of the severity of winter weather for maintenance operations on Provincial highways". Synectics has been working with researchers at the University of Waterloo, AIRD of Environment Canada, IMOS Inc. and ASI Technologies.

To date, the research team has reported on:

- The appropriateness and quality of winter maintenance data for Ontario Provincial highways, and in particular the promising nature of Automated Vehicle Location (AVL) data and the need to do data quality checks on Maintenance Management Information System (MMIS) data;

- The results of multiple regression models using daily salt use (tonnes per lane-kilometre) as the dependent variable, focusing specifically on selected routes in the Chatham, New Liskeard, Ottawa and Thunder Bay-Kenora areas;

- The limited value of incorporating terrain variables related to drifting such as "percent of the highway segment that is wooded" in modeling daily patrol-level maintenance activity;

- The challenges associated with relying exclusively on one source for weather data;

- The need to incorporate either data filters or classification techniques into the development of winter indices. The reason for this is that individual weather variables have different effects, depending on their co-occurrence or the sequence of events. For example, current models do not adequately predict maintenance activity for ground frosts, early winter snowfalls, some freezing rain events, extended periods of very low temperatures, and prolonged winter storms where operators are working

In March, 2007 Synectics submitted its recommended next steps in the project, which included a more detailed analysis of how decisions are made. The current project addresses this recommendation by focusing on a smaller unit of analysis—the individual truck on an hour-by-hour basis—and modeling these data in non-traditional ways that incorporate a wide variety of weather data. In the current project, salt is applied on each individual highway segment by only one truck; thus the results can be interpreted as being either truck-specific or highway segment-specific, since the two are interchangeable in the case study area.

OBJECTIVES

The overall goal of the project is to build on the work that has been completed on winter indices by the Synectics team by using classification techniques of various types. The specific objectives are as follows:

a) To use reliable data at a fine spatial-temporal resolution to characterize both weather conditions and winter maintenance decisions. The project focused specifically on the decision to apply rock salt at particular rates.

b) To explore the value of different types of classification modeling in linking salt application decisions to winter weather variables—both forecast and observed.

c) To explore, to the extent possible given the data, the effects of direct liquid application on road salt usage.

Approach and data

GENERAL ANALYTIC FRAMEWORK

In any modeling exercise, it is important to begin with a conceptualization of what it is that we want to understand. In modeling winter maintenance, the dependent variable is usually some measure of resource use, which is based on human responses to a combination of expected and observed physical conditions. As a result:

1. Human behaviour in the form of operation decisions is, in essence, the process that we are trying to understand and/or prescribe in models of winter road maintenance. Accordingly, any model of past activity will likely have a local signature, i.e., it will incorporate local knowledge and will reflect local practices. Since it is reasonable to assume that different combinations of activities may have similar effects with similar costs (Mahoney et al., 2005), it is unrealistic to expect that identical weather and road conditions will result in identical treatment on an hour-by-hour and truck-by-truck basis. Thus as a starting point, we do not expect any model of multiple jurisdictions or highway segments to achieve 100% fit based on weather data alone. This has implications for the generalizability of model specifications—an issue that requires more attention than has been previously acknowledged.

2. Operational decisions (and adjustments to those decisions) are made continually at the truck or patrol level; discerning the effect of such decisions requires data in small units of time and space. Therefore, even in instances where fine-scale predictions do not produce excellent fit or where only aggregate results are required for decision making, it is important to work with a small spatial-temporal unit of analysis if error terms are to be understood, and models are to be improved.

3. Although maintenance activities, when aggregated over time and space, define continuous variables (e.g., weekly salt totals, seasonal equipment hours), individual decisions are usually categorical. Indeed, the decision to treat is binary, and there is a finite number of discrete treatment modes used by any one truck on any single highway segment. It is therefore important to explore modeling approaches than are consistent with the inherent nature of the dependent variable(s).

Despite the importance of the above issues, most models reported in the literature have over-simplified the characterization of decision making by maintenance personnel, have modeled maintenance activity in coarse spatial-temporal units, and have treated the dependent variable as continuous. Examples of such models include the Hulme (Hulme, 1982; Cornford, D. and Thornes, 1996) and SHRP (Thornes, 1993; Andrey et al., 2001; Decker et al, 2001) indices.

In terms of the modeling approach, some studies begin with no prior assumptions about the importance of specific weather variables (Suggett, 2007), while others attempt to capture different types of treatable events; an example of the latter is the inclusion of snowfall and frost events as additive terms in the SHRP model (Thornes, 1993). Regardless, most models are based on some type of regression modeling, and thus are constrained by the issue of multi-collinearity and the limitations of linearly additive models. Alternate approaches, such as storm event classifications (Nixon and Qiu, 2005) are also being explored, but their association with specific treatments has yet to be objectively validated using recorded maintenance data.

In terms of model fit, some past studies have report r-squared values greater than 0.9 (Thornes, 1993, Boselly et al, 1993), but these are based on highly aggregated data. In addition, when transferred to other regions, these models do not perform well, even when re-calibrated locally (Andrey et al., 2001; Decker et al., 2001). Models that have been developed at somewhat finer spatial-temporal units (Carmichael et al., 2004) typically have achieved much poorer fit, also limiting their usefulness. There is the added problem that some key variables produce different coefficients in different circumstances, and in virtually all cases assumptions about data distributions, linearity and independence are violated. The net effect is that many winter severity models and indices are useful only for benchmarking use of materials from season to season in the same geographic area. They add little to our understanding of what triggers or sustains specific maintenance operations.

The current project introduces a new approach for modeling winter maintenance using classification techniques applied to hourly level data for individual highway segments. The dependent variable is the salt treatment mode, as explained in more detail below, and the independent variables include numerous forecast and observed weather variables.

STUDY AREA

The current project uses data from Highways 417 and 138 in the Province of Ontario, Canada. These highways are located in the VanKleek Hill patrol of the Ottawa District. Highway 417 is a four-lane freeway designated by the Ministry of Transportation as Winter Service Class 1, running 135 km from Ottawa, Ontario toward Montreal, Quebec; whereas Highway 138 is a 35 km, Class 2, 2-lane rural highway. The VanKleek Hill patrol contains approximately110 km of Highway 417 and all of Highway 138.

Winter Service Class 1 stipulates that snow removal begins before 2.0 cm of snow accumulation on the roadway, plowing and salting continues with 1.3 hour circuit time, and pavement is essentially bare within 8 hours of the end of the storm. Service Class 2 stipulates a 1.8 hour circuit time and that pavement is essentially bare within 16 hours. Winter sand is applied when temperatures fall below -12˚C and snow cannot be removed by plowing and salting. In addition, anti-icing liquid is applied by truck to bridge decks and other ice-prone spots on Hwy 417 in advance of storms, except at one location where it is applied by a Fixed Automated Spray Technology (FAST). Granular salt is pre-wetted with anti-icing liquid and applied at rates varying with pavement temperature and snowfall, from 50 to 170 kg/2-lane km (6.8-23.3 g/m^2). Plows and combination plow-spreader units are used to anti-ice and clear snow. Winter operations are planned using data from the provincial Road Weather Information System (RWIS).

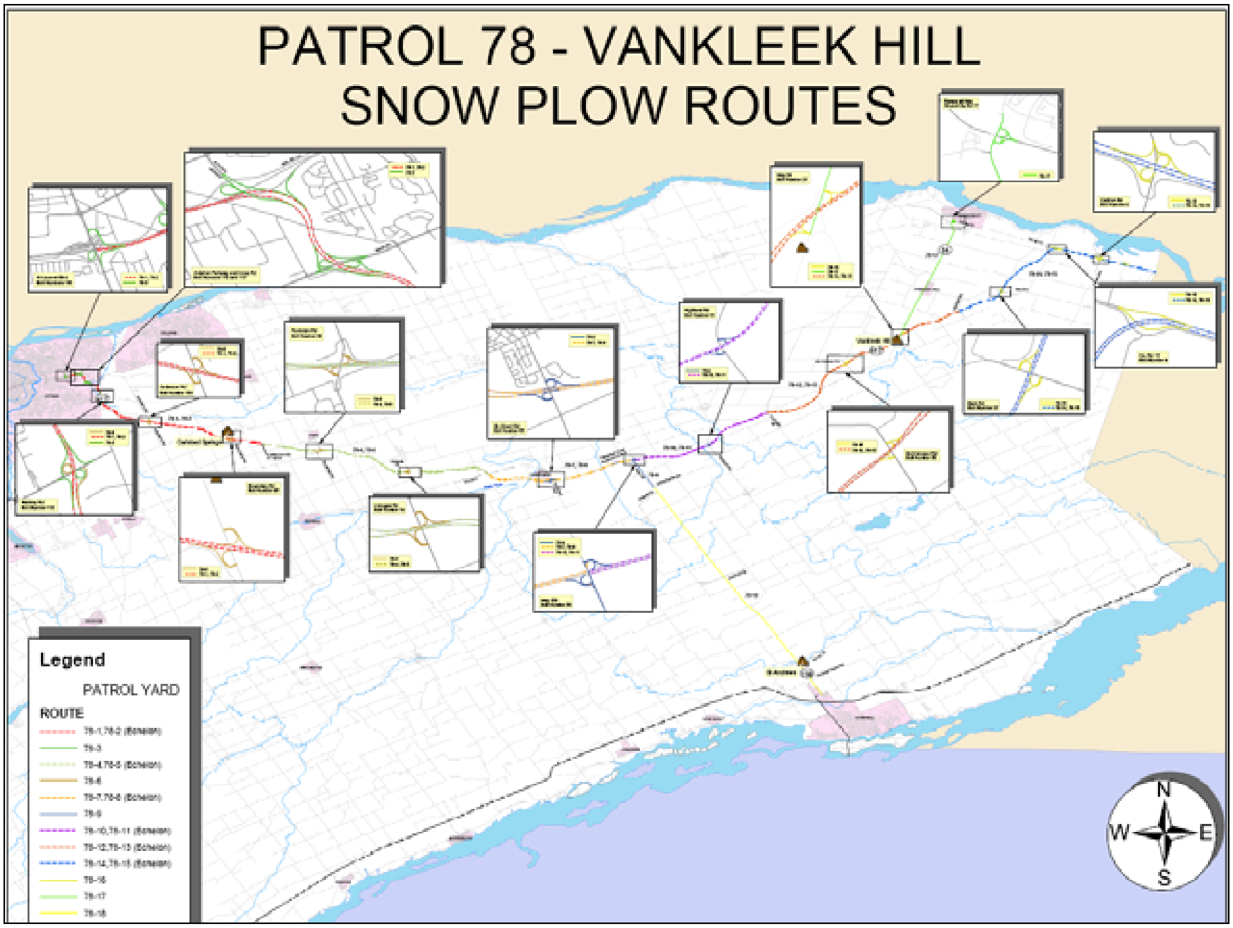

Figure 1 – The study area including information on the individual trucks/highway segments

AVL DATA AND THE DEPENDENT VARIABLE

Winter maintenance activity data were extracted from the internet-based system of Automated Vehicle Location (AVL) records. These were extracted using the "Query" and "Winter operation report" (WOR) functions, and were downloaded as CSV spreadsheets. The query function, which provides readings in time intervals as short as 5 seconds, was used to extract details on maintenance operations. Each reading provides information on materials used, driving speed and exact location. Data were extracted for three seasons, from November1 to April 30, 2004-2007.

Winter operation reports were used to validate daily and seasonal materials usage. They provide information on trips, i.e., vehicles entering and leaving yards (Ministry of Transportation, Ontario), and include the date of the report as well as the start time and return-to-yard time. They also include total kilometers serviced, material usage and average rate of application during the treatment period.

The unit of analysis adopted for the study is the truck-hour. A spreadsheet was developed where each row represents a unique hour for one truck and the columns include information on the dependent variable (treatment activity) and the independent variables (primarily weather variables). To simplify the task, we focused only on material application rates and used these to create what we called "treatment modes" (e.g., a treatment mode of 65-39 means that the salt application rate was 65 kg/ km and the pre-wet application rate was 39 liters/tonne of salt).

Since the data set used for modeling contains information on every hour from November 1st to April 30th, the AVL system was not active for a majority of the hours; these hours were coded as n/a. The number of n/a readings varied by truck from 10,892 to 12,466, out of the total of 13,032 hours during the three-season study period. In total, 86 percent of hourly records were n/a hours indicating no AVL readings. During the remaining 14 percent of the hours, when there were AVL readings, zero was the most common reading, occurring between 34 and 69 percent of the time, depending on the truck and season. In these instances, the truck was active but salt was not being applied. The focus of the modeling exercise was to separate out these n/a and zero hours from the non-zero or "treatment" hours (approximately 7 percent of the total) based on forecast and observed weather.

Preliminary examination of the AVL data indicated a small number of treatment modes. The dominant salt application rates were 100, 130, and 170 kg per two-lane km (as well as 50, 65 and 85 mostly for one-lane treatments on Highway 138). There were also relatively frequent readings of 70 and 260 kg/km travelled. The frequency distributions for the seven trucks (7805, 7808, 7809, 7811, 7813, 7815, 7818) used in the analysis are provided in Figure 2. It is worth noting that the instantaneous AVL readings showed a fair degree of oscillation around these values, partly due to mechanical lag as the spreader attempts to maintain a uniform application rate in response to varying the truck's speed, but also because treatment was sometimes intermittent or changed within the hour. Individual readings were thus assigned to the closest of the primary values noted above.

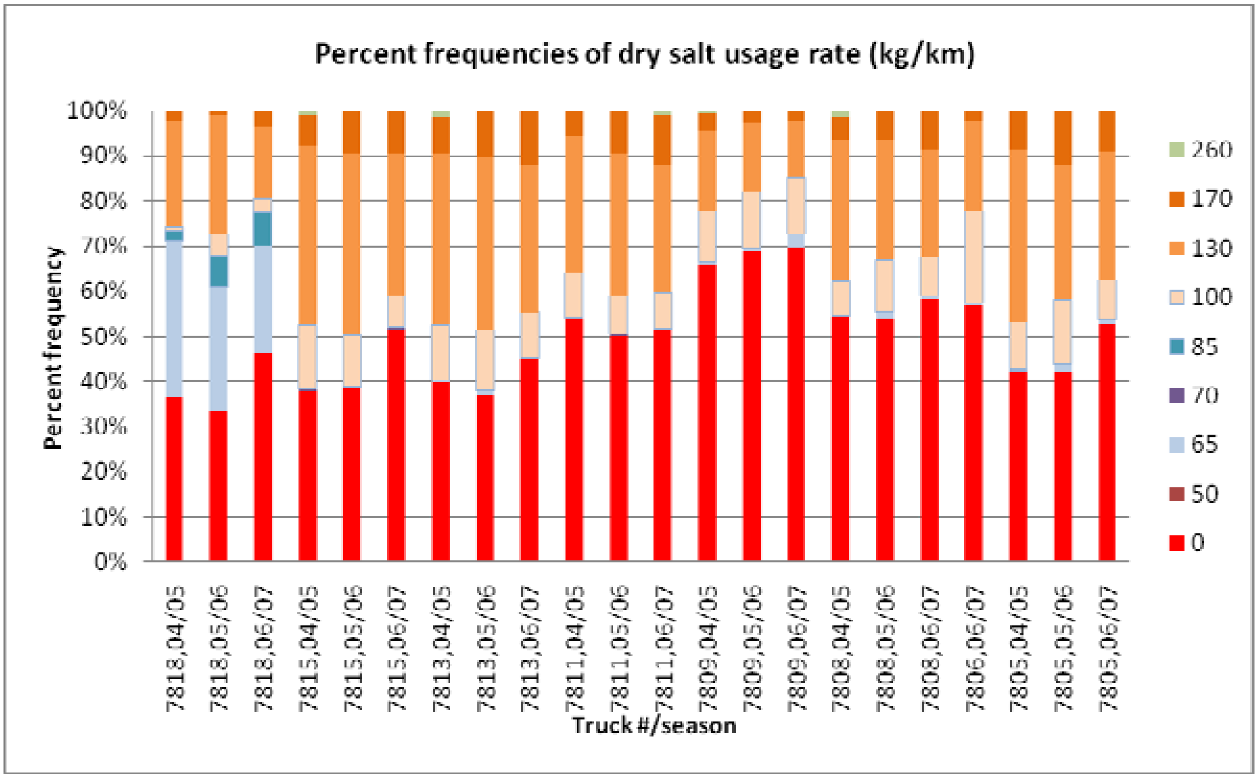

Figure 2 – Salt application rates by truck and season

The first step taken in creating the dependent variable was to aggregate the AVL readings for an individual truck for each hour using pivot tables so that the frequency distribution of treatment modes, i.e., application rates of salt and pre-wet, could be examined. An hour in which the truck was active for the whole 60 minutes produced up to 300 readings, but the number varied depending on how the truck was operating, i.e. spreading salt or just plowing. This pivot table was then used to determine the "dominant mode" for that hour. More specifically, if during a one-hour period the vehicle had more that 75 percent of the readings recorded as zero, then that hour was assigned a value of zero, indicating that no dry salt was applied. If material application was observed for at least 25 percent of the readings in an hour, then the mode that had the most readings was chosen as the dominant mode for that hour. Sensitivity analysis was performed with different cut-off values and assignment rules, but the above worked comparatively well in reproducing daily and seasonal salt usage.

Second, the hourly treatment modes were combined with average truck speed and the percentage of the time that treatment occurred in “treatment hours” to estimate daily and seasonal salt use. Figure 3 presents a comparison between the daily estimate and the actual amount of salt used for one of the trucks for one season. The estimate of dry-salt usage was within 5 percent of actual usage. On a daily basis, the fit was within 1 tonne for most days. The other seven trucks typically performed better than the one illustrated. This suggests that if the modeling exercise can accurately assign hours to a treatment mode, then it is possible to estimate seasonal salt usage with a high degree of accuracy and daily or storm-level salt usage with a fair degree of accuracy.

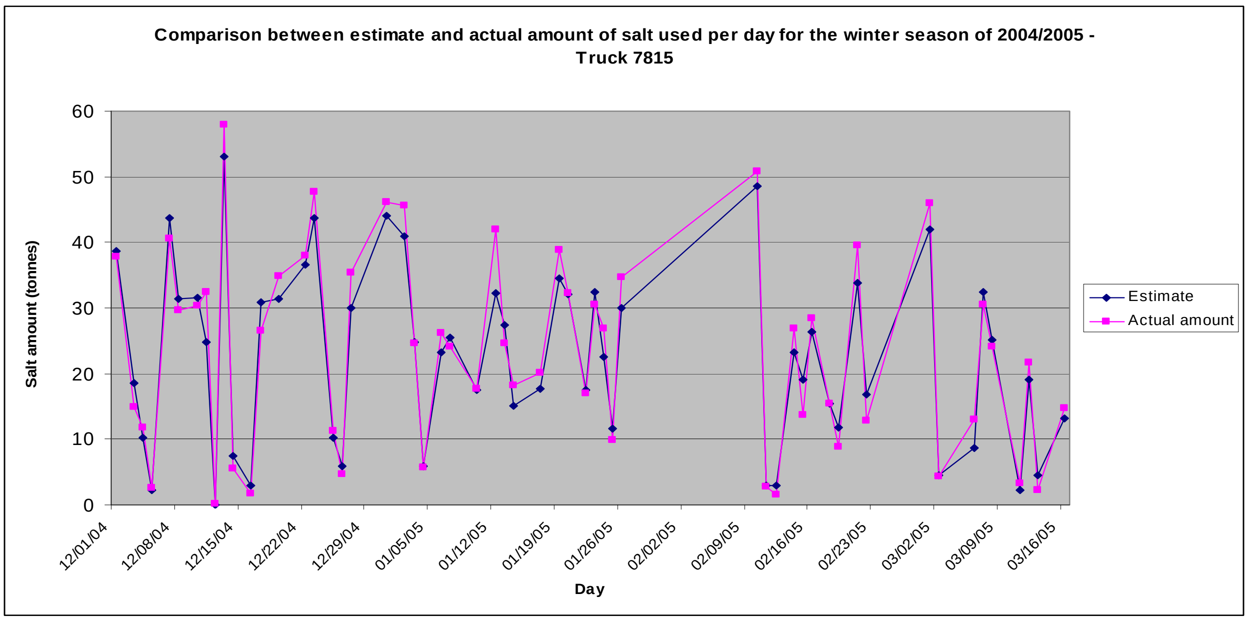

Figure 3 – Comparison between daily estimate and actual salt use (Truck 7815)

Once the treatment modes were determined for individual truck-hours, these data were examined in terms of their temporal variations. The data make it clear that treatment patterns are not the same for all trucks operating in a patrol, highlighting the importance of highway-specific variables and the limitations of using only weather data to model and ultimately predict salt usage. Table 1, which provides salt data for February 14, 2007 for four trucks, illustrates this point. Using the approach outlined in the previous paragraph, these treatments would have resulted in the following daily totals in salt usage (kg) for trucks 7805, 7808, 7811 and 7813, respectively: 58, 42, 42 and 60 tonnes.

While these values are similar, they are not identical even though they are in response to the same general weather conditions. Unless, highly relevant and site-specific variables can be included in future work, there will always be some unexplained variation in salt use.

Table 1 – Example of how salt application rates (kg/km) vary

Hour	Truck 7805	Truck 7808	Truck 7811	Truck 7813
00	0	0	0	0
01	0	0	0	0
02	0	0	0	0
03	0	0	0	0
04	0	0	0	130
05	130	130	130	130
06	170	130	130	130
07	130	130	130	130
08	130	130	0	130
09	170	170	130	170
10	170	170	170	170
11	130	0	170	170
12	130	0	170	170
13	130	170	130	130
14	130	130	0	130
15	130	130	130	130
16	130	170	130	170
17	170	170	0	130
18	170	0	170	130
19	130	0	170	130
20	130	130	130	130
21	100	130	0	0
22	100	0	0	130
23	130	0	0	130

Also, even within treatment periods when there is little variation in weather, material application rates vary, in response to past treatment and the observed road conditions. This is illustrated in Figure 4 for truck 7818 on Highway138 for over a two-day period in March, 2005. Truck 7818 spot treated certain stretches of road on the morning of March 6th and then applied salt on a nearly continuous basis, at variable rates, for more than 24 hours. Light snowfall occurred on the morning of the 6th and somewhat heavier accumulation occurred over a 12-hour period starting on the morning of the 7th.

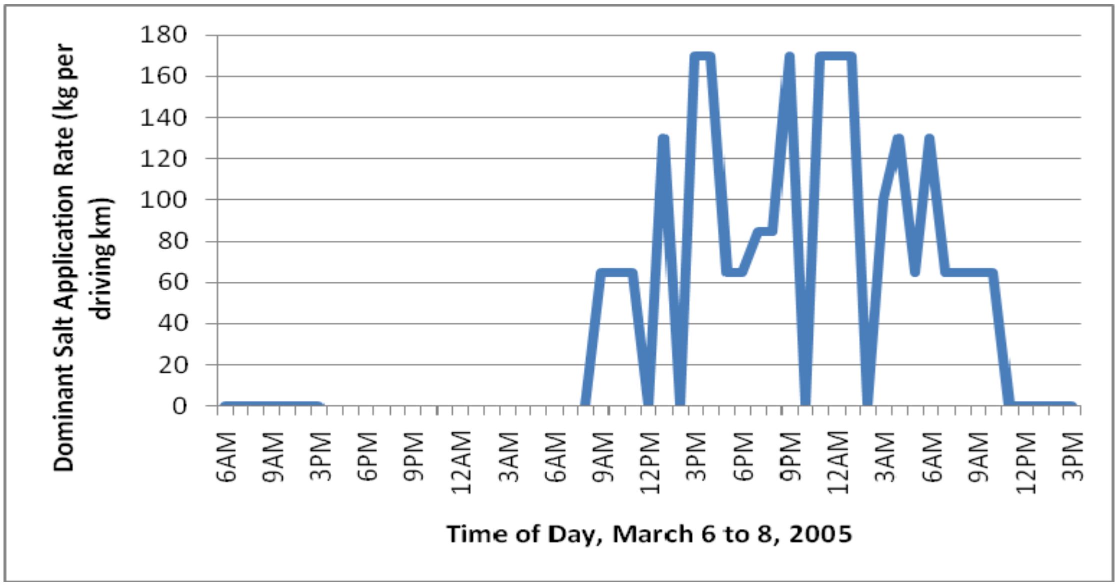

Figure 4 – Temporal Variation in Salt Application Rate (Truck 7818)

The number of salt application rates was still fairly large after the assignment procedure was applied such that it would complicate the classification modeling. We therefore further simplified the dependent variable as follows. Each hour was assigned to one of the following categories of dry salt usage: n/a (no AVL data for this hour indicating that it is not in a treatment period), zero (truck controller is active but no/little dry salt was applied, for example when driving back to yard, plowing only or doing spot treatment), low (~50 kg of salt per lane-km, i.e. controller readings of 50 or 100), medium (~65 kg of salt per lane-km, i.e. controller readings of 65, 70, or 130) and high (~85 kg of salt per lane-km, i.e. controller readings of 85, 170 or 260). From the point of view of salt usage, the n/a and zero categories are equivalent. The reason for treating them separately is that we would expect different weather conditions to be associated with each: the n/a would be associated with times when snowfall, icy conditions, and drifting are all absent; by contrast the zero category would reflect winter weather conditions when dry salt was not being continuously applied but maintenance activities were being performed (e.g., plowing during very cold weather).

EXPLANATORY VARIABLES

A summary of the explanatory variables incorporated in the analysis is presented in Table 2, and a full listing in provided in the Appendix. The selection of variables was driven by two primary assumptions: 1) that winter maintenance operations are influenced by both the expected and the actual occurrence of particular weather and road conditions; and 2) the frequency, timing, intensity and duration of these conditions is related, though not necessarily directly/linearly, to the amount of salt applied. Thus variables describing both forecast and observed conditions were chosen. In general, the location and spatial/temporal resolution of the dependent variable dictated the form of the explanatory variables.

Table 2 — General description of explanatory variables

General category	Temporal resolution	Variable description
Weather Forecasts	Semi-daily (tonight 1800-0500/tomorrow 0600-1700)	occurrence, type, amount and probability of precipitation occurrence of blowing snow temperature and trend
Weather Warnings	Hourly	watches, warnings and advisories for winter storm, heavy snowfall, freezing rain, flash freeze, or blowing snow conditions
Standard Weather Observations	Hourly	visibility, wind speed and direction, air and dewpoint temperature occurrence, type and intensity of precipitation occurrence of blowing snow
	6-hourly	liquid-equivalent precipitation accumulation
	Daily	snowfall, rainfall and total liquid-equivalent precipitation accumulation minimum, maximum and mean temperature depth of snow on the ground
Radar-derived Observations	Hourly	presence of precipitation near or over study area estimated average and peak precipitation intensity over study area
Road Weather Conditions	Hourly	air and dewpoint temperature occurrence, type and intensity of precipitation pavement surface temperature road surface condition (e.g., dry, wet, icy, etc.)

Weather forecast and warning data were obtained from Environment Canada for the Prescott and

Russell forecast region in which the study area is located. Short-term weather forecasts issued in the afternoon of a given day predict conditions for that night and the following day. For the analysis, we have assumed that tonight conditions are valid from 1800-0500 hours and tomorrow conditions represent 0600-1700 hours. Inclusion of forecast variables is an explicit attempt to recognize the decision-making behavior of maintenance staff in advance of sensible weather or road conditions. Like forecasts, the issuance of weather warnings (watches, warnings or advisories) is assumed to provide maintenance personnel with early notice of treatable events and greater confidence that they will occur.

As noted, past research efforts have often relied exclusively on standard weather observation data to model road salt and other measures of winter road maintenance. Although the current study incorporates other sources of information, standard observations are also an important set of explanatory variables. Data for the variables identified in Table 2 were obtained from Environment Canada for the Ottawa MacDonald-Cartier International Airport observing station. While this is the closest primary observation station—and thus collects the full suite of required variables with high quality and continuity—it is situated in the extreme northwest corner of the study area and may not be fully representative of conditions, especially for precipitation.

Inclusion of radar-derived estimates of precipitation occurrence and intensity should help to address inadequacies in standard observations related to areal coverage. Hourly radar imagery (CAPPI, snowfall mode) was obtained through Environment Canada for the Franktown radar station located about 80 km west of the study area. For the pilot study, images were subjectively interpreted to determine the location and intensity of precipitation.

Road Weather Information Systems (RWIS) can also be used to improve the spatial representativeness of standard observing stations and, more importantly, provide several variables unique to the road condition. Hourly measurements were selected or derived from data for the Casselman RWIS station (Eastern Region-16) obtained from the Ontario Ministry of Transportation.

For this exploratory work, highway geometrics and other road attribute data were not included.

However, potential influences related to the different highway segments and to seasonal and diurnal variations were represented by truck, month and hour variables.

Modeling

Two complementary modeling approaches were used in exploring the relationships between weather and salt application rates on Ontario highways— classification trees on their own, and then two types of direct and indirect models, three of which include tree-based techniques. All analyses were conducted in the data analysis environment R (Ihaka, and Gentleman, 1996) and its packages, "rpart", "ipred", and "mda". In all cases, the modeling was completed using data from three seasons (13,032 hours) and seven trucks from the Vankleek Hill Patrol near Ottawa, Ontario.

CLASSIFICATION TREES

A classifier is a function that predicts the (unknown) class membership of an object based on (known) explanatory variables. In supervised classification, as was done in this study, this function is fitted to a training data set for which the class membership of each object is known. The fit of the model can then be tested on a different test data set, to avoid overfitting and reporting over-optimistic error estimates. In the present setting, the objective was to model the ordinal-scale treatment mode of truck-hours as a function of the weather variables presented above. To account for the potential delay between treatment decision and implementation, variables shifted back by one hour and two hours were also considered. Future work will incorporate a 24-hour lag as well, to account for the clean-up effects of major snowfalls as recorded in the daily records of snowfall amount.

In this exploratory analysis of winter maintenance activities we used classification trees because they provide a simple means for modeling complex interactions between explanatory variables (Breiman et al., 1984). This appears to be more appropriate in this context than the additive combination of variables by linear or logistic models.

Classification trees are hierarchical sets of yes/no questions that are derived from a training data set by minimizing a loss function. In statistics, a loss function represents the cost associated with being wrong in certain ways. For instance, this function may allocate different costs to, for example, false-positive predictions (incorrect prediction of belonging to a treatment period) or the misclassification of a high-treatment truck-hour as a low-treatment hour. We use pruning to cut down the classification tree to a reasonable size and avoid over-fitting the classifier to the data.

INDIRECT MODELS

Indirect models try to make use of intermediate variables to improve the prediction of structured data. Examples are structured class definitions in classification (Hand et al., 2001; Peters et al., 2005), or a priori knowledge on the structure of regression relationships ("coaching variables" in regression; Tibshirani and Hinton, 1998). In conventional direct modeling, by contrast, a response variable y is predicted directly from some combination of predictor variables $x_1,\ldots,x_n$. In many practical situations such as medical diagnostics, however, there are additional variables $w_1,\ldots,w_k$ that are more closely related to the response, but that are expensive or hard to obtain (e.g. from tissue samples or specialized technology) than the predictors x_i, which are available in the standard situation. If the variables w_i are known on some learning sample, then it is often possible to obtain improved indirect predictors of y by first predicting the so-called intermediate variables w_i from the predictors x_i, and in a second step the response y from the predicted intermediate variables.

The previous discussion of the dependent variable illustrated that accurate predictions of hourly treatment modes will result in accurate estimates of aggregate salt use. The key therefore is to build an indirect model that predicts treatment modes from weather and road conditions, and then integrate these intermediate predictions into the estimation of salt use. Thus, the final predictive model only depends on the predictors $x_1,\ldots,x_n$, but exploits our knowledge of the structured relationship between salt use and treatment mode.

The concept of indirect modeling was introduced by Hand et al. (2001) and extended by Peters et al. (2005) in the context of classification using machine-learning techniques, especially ensemble methods. Tibshirani and Hinton's study (1998) on coaching variables in regression shows that a similar approach can be useful in regression. We will therefore extend the generalized indirect classification approach of Peters et al. (2005) to a regression context where classification techniques serve as intermediate models. We compare and combine both machine-learning techniques and statistical techniques in two indirect models, and compare their results with two direct models:

- Linear regression model with stepwise forward variable selection (acronym: LM): This is a standard statistical technique and the most commonly used form of direct modeling. In the current analysis, the number of variable selection steps (based on the Akaike Information Criterion, AIC) was limited to 30. Linear models may give good explanatory models when few variables are involved, and are relatively resistant to overfitting because of their inflexible linear structure. The coefficients may be misleading, however, because of problems of non-linearity or multi-collinearity.

- Bagging: Bootstrap-aggregated regression trees, which are another form of direct modeling, were obtained by averaging over 50 individual regression trees that were grown up to a maximum depth of 15. This results in a highly flexible regression technique that should be able to detect higher-order interactions between predictor variables, such as influences of certain weather variables only during specific storm types. Bagging was introduced by Breiman (1996).

- Linear regression combined with bundling (acronym LM-bund): Here the bundling approach of Hothorn & Lausen (2005) was used to build an efficient classifier that combines the capabilities of a tree-based classifier (as in bagging) with the advantages of an ancillary classifier (here: penalized linear discriminant analysis, PLDA; Hastie et al., 1995) capable of detecting linear class boundaries. This is an example of indirect modeling in that both the predicted intermediate variables and the original predictors were fed into a linear regression with stepwise forward variables selection (maximum 30 steps) to obtain results that were compared to the direct LM approach. We used 25 trees with a maximum depth of 7 for bundling.

- Bagging with regression trees and PLDA (Bagging-CT/PLDA): This is an indirect version of the direct bagging technique, in which two classification techniques (individual classification trees, Breiman et al., 1984; and PLDA) were used as competing intermediate classifiers, each predicting treatment modes and their probabilities. The simpler CT/PLDA was preferred to a nested bundling classifier to avoid a prohibitive increase in computing time. Pruning was applied to avoid overfitting CT growth.

In building the indirect models, the training dataset was subsampled (on the day level) to obtain two independent data sets of equal size; the intermediate classifiers are trained on one of these sets, and the regression technique on the other one in order to avoid overfitting (compare Peters et al., 2005).

We used weather variables, as listed in Appendix A, as predictor variables (in addition, effects related to truck, month, and hour of the day were allowed), the daily amount of salt use (average amount in kg/h) as response, and the predicted treatment modes (n/a, zero, low, medium and high) and number of lanes treated (n/a, zero, one and two) and respective probability predictions as intermediate variables.

An important issue in the comparison of different predictors is the unbiased error estimation. Modern resampling-based approaches are the cross-validation and the bootstrap (e.g., Efron and Tibshirani, 1993; Hand, 1997). The bootstrap consists of resampling of observations with replacement, thus simulating the generation of a new random sample from the unknown random process that is approximated by the empirical distribution of the observed data set. Two independently sampled bootstrap samples from the available data are then used as training and test data sets. The sub-sampling is carried out at the day level because different hours of a day are likely to show very similar weather and road management conditions. Note that this grouping by day also implies that data from different trucks on the same day is kept together in resampling. Similar approaches for resampling dependent or grouped data have been applied elsewhere (compare Bühlmann, 2002;

Davison et al., 2003; Brenning, 2005; Brenning et al., 2006; Brenning and Lausen, 2008). We report the average accuracies (r-squared) of the models estimated on the test data sets.

While the bootstrap is usually based on randomly drawing N out of the N available samples, we only use bootstrap samples of size $0.4N$ because of the large size of our data set. A favorable side effect of sub-sampling is that the number of objects (here: days) sampled into both the training and test set is greatly reduced, making a bias reduction (such as the .632 bootstrap) unnecessary.

Results

CLASSIFICATION TREES

In this exploratory analysis of winter maintenance activities we used classification trees to distinguish between five salt application rates (n/a, zero, low, medium and high), based on AVL hourly data. The classification tree algorithm was run for each of the seven trucks separately and then for all seven trucks together. Trees were grown based on the training data set and then tested on a second, independently drawn sub-sample (drawn on the day level) of the same size. In growing the trees, the cost structure reflected the importance of non-zero treatment modes, i.e., it put little penalty on misclassifying an n/a as zero or vice versa, moderate penalty on incorrectly assigning a low, medium or high treatment modes to another one of these classes, and high penalty on assigning an n/a or zero to a non-zero mode or vice versa.

The results, which are summarized in Tables 3 and 4, provide an indication of the overall high level of accuracy achieved in assigning treatment hours to the correct category. For all trucks together, 81.3 percent of all hours were correctly assigned and for individual trucks the percentages varied from 81.1 to 86.1. As well, the models performed well in correctly assigning hours where no salt was applied to either the n/a or zero class. However, the models performed less well in designating treatment hours as such, which would overall result in an underestimation of aggregate salt use. Future work will concentrate on adjustments to the loss function in order to improve the assignment outcomes. Another approach for dealing with the underestimation is to use the tree-based assignments along with weather variables to derive an unbiased indirect model of salt use, as discussed in the next section.

Table 3 — Classification Tree Results, All Trucks Together

		PREDICTED					
		n/a	zero	low	medium	high	# hours
OBSERVED	n/a	**27627**	2112	0	849	195	30783
	zero	1194	**590**	0	799	69	2652
	low	128	99	**0**	299	9	535
	medium	303	217	0	**1086**	39	1645
	high	71	62	0	299	**3**	435

Accuracy (% hours correctly assigned) =	81.3
Specificity (% hours correctly assigned to n/a , zero) =	94.3
Sensitivity (% hours correctly assigned to non-zero) =	66.3

Table 4 — Classification Tree Results, Individual Trucks

	TRUCK						
	7508	7808	7809	7811	7813	7815	7818
Accuracy (% hours correctly assigned)	81.7	84.3	86.1	83.7	85.6	81.1	82.2
Specificity (% hours correctly assigned to n/a , zero)	95.2	97.1	98	95.1	93.1	94.3	94.3
Sensitivity (% hours correctly assigned to non-zero)	66.5	58.1	33.6	69.3	77.4	70	66.1

Before turning to the preliminary results on indirect modeling, however, it is important to comment on the nature of derived trees, since this is perhaps the most promising aspect of this approach for potentially building decision-support systems. The promise lies in the fact that classification trees identified both variables and threshold values that are consistent with expert knowledge and scientific principles. While individual trees differ, and thus we are not yet at a point where we can identify a robust tree with a high degree of fit, the results for pruned trees, based on a 40 percent sample of data from all trucks, illustrate the results.

Table 5 summarizes the two paths that predict high salt application rates—typically 170 kg/two-lane km. The first tree path describes those conditions when then is little snowfall but freezing rain is expected or occurring, and the second situations depicts snowfall occurring with mild temperatures. In both instances, a mix of information sources (e.g., standard weather observations, RWIS, radar) and information at various time stamps (e.g., recent past, present, forecast) is used to define the trees, similar to what would occur in the yard.

This mixing of information types is also evident in the three tree paths that predict medium salt application rates—typically 130 kg/two-lane km. The first of these three paths describe conditions when there is light rain, freezing rain or snowfall and pavement temperature below 1 degree Celsius. The second and third paths describe situations when snowfall is occurring; for the second path the snowfall is widespread and there are gentle or moderate winds, and for the third there is lower visibility and the pavement surface is below zero.

Table 5 — Tree Paths Predicting High Hourly Salt Application

	1st tree to predict high salt application	**2nd tree to predict high salt application**
1st split	Daily snowfall amount, 2 hrs ago, < 0.7 cm	Daily snowfall amount, 2 hrs ago ≥ 0.7 cm
2nd split	Slight snow or unidentified precipitation (RWIS)	Visibility distance ≤ 14.5 km
3rd split	Daily snowfall amount, 2 hrs ago, < 0.025 cm	Pavement surface temperature < 0.15°C (RWIS)
4th split	Freezing rain warning, 2 hrs ago	Daily snowfall amount ≥ 2.1 cm
5th split	Pavement surface temperature, 2 hrs ago, < 0.9°C (RWIS)	Air temperature, 2 hrs ago, ≥ 0.5°C
6th split	Dewpoint temperature, ≥ -6.6°C (RWIS)	

Table 6 — Tree Paths Predicting Medium Hourly Salt Application

	1st tree to predict medium salt application	**2nd tree to predict medium salt application**	**3rd tree to predict medium salt application**
1st split	Daily snowfall amount, 2 hrs ago < 0.7 cm	Daily snowfall amount, 2 hrs ago ≥ 0.7 cm	Daily snowfall amount, 2 hrs ago ≥ 0.7 cm
2nd split	Rain, snow or freezing rain (RWIS)	Visibility distance >14.5 km	Visibility distance ≤ 14.5 km
3rd split	Pavement surface temperature, 2 hrs ago, ≤ 0.9°C (RWIS)	Precipitation echoes cover at least 25% of the study region (Radar)	Pavement surface temperature, < 0.15°C (RWIS)
4th split		Wind speed at 10 m ≥ 12 kph	

It is worth noting that, despite the interpretability of these tree paths, the model did not perform well at predicting hours of high salt application (Table 3), with many false positives and false negatives. For moderate salt application rates, the model performed better with 1086 correctly assigned hours, but there were also a large number of hours where medium-level salt application occurred but the model predicted no treatment (520 hours) or a high rate of salt application (39). There were also numerous hours assigned to the medium treatment category, when either no treatment was observed (2246 hours) or where a different level of treatment was observed (299 high and 299 low). Future work will focus on improving the performance of tree classifiers, as discussed later.

INDIRECT MODELS

In this section, we discuss the results of four different modeling approaches—two direct and two indirect. The goal of this exercise is to demonstrate the improved model fit of newer, more sophisticated approaches relative to linear regression, which is the basis of many previous models. In order to achieve robust results, each of the four models were run multiple times—initially 14 times (Tables 3 and 4 below reflect the values for these), and then later for a total of 100 times (with nearly identical results to what was achieved with the initial 14). For each of the "runs", the models were calibrated on a training data set that comprised a randomly selected subset of 40 percent of the days in the data set. The estimated values of salt use were then calculated for a test data set, comprising a 40 percent sample, and aggregated to the day. The latter were then compared to the salt usage that would have occurred had each of the 24 hours been perfectly assigned to their respective salt application rates on that day.

Results are summarized in Table 7, and results of Wilcoxon signed rank tests for pairwise comparisons of the r-squared estimates are given in Table 8. The results show that the indirect method of bagging with intermediate classification tree and PLDA models (Bagging-CT/PLDA) performs significantly better than the other three approaches, yielding a bootstrapped r-squared value that is between 6.1 and 8.1 percentage points higher than was found for the other methods. Pairwise differences between the other models are between 0.1 and 2.1 percentage points; however they are not significant at the nominal 5 percent significance level from each other with only 14 bootstrap replications. (Additional bootstrap replications that were later computed confirm the small magnitude of these pairwise differences, but provide significant results because of the larger number of replications.)

Table 7 – Estimates of R^2 based on 14 replications

	Method	R^2 (std. dev.)
Direct methods	LM	61.6% (4.3%)
	Bagging	63.6% (7.4%)
Indirect methods	LM-bund	61.5% (5.0%)
	Bagging-CT/PLDA	69.7% (5.0%)

Table 8: Mean differences in estimates of R^2 and results of Wilcoxon signed rank tests at the 5% significance level

	Direct Models		Indirect Models	
	LM	Bagging	LM-bund	Bagging-CT/PLDA
LM		-2.1%	-0.1%	**8.1%** ***
Bagging	2.1%		-2.1%	**6.0%** **
LM-bund	0.1%	2.1%		**8.2%** ***
Bagging-CT/PLDA	-8.1%	-6.0%	-8.2%	

(*** *p*-value <0.001; ** <0.01; * <0.05).

The indirect modeling approach that combines tree-based and linear intermediate learners of treatment mode with a tree-based ensemble learner of the final response variable (salt use) achieved statistically significantly better results compared to direct approaches and compared to an indirect approach that assumes a linear relationship between predictors and salt use response. The improvement achieved with this novel modeling approach compared to more traditional ones is not only statistically significant, but also of a magnitude (6.1-8.1 percentage points of the R^2) that is substantial from a practical perspective. The use of bootstrap estimation techniques ensures that these error estimates are unaffected by potential overfitting of models.

In addition to the superior predictive performance of the indirect "Bagging-CT/PLDA" approach, it has to be pointed out that this "black-box" machine-learning approach is not less attractive in terms of model interpretability than the direct linear regression model with stepwise variable selection from the available predictors. While the linear regression model is often assumed to produce an interpretable "white-box" model with a simple model structure and simple relationships between individual predictors and the response, the result of a stepwise variable selection with up to 30 predictor variables in the final model (as in the present study) can, if anything, result in misleading interpretations of the model structure. The reason for this is in highly correlated variables with very different semantics, nonlinear relationships that are captured by combining several correlated variables with counterintuitive coefficient estimates, and unobserved confounders. Thus, model interpretability is not lost by using indirect and/or machine-learning techniques. Moreover, the indirect approach is able to capture our knowledge about the structured relationships between weather and salt use amounts through the selection of treatment modes (compare Hand et al., 2001; Peters et al., 2005), resulting in a "modular" model design and the possibility of evaluating model components such as the intermediate models separately.

Conclusions and Recommendations

CONCLUSIONS

Models of winter maintenance activities, and associated winter weather indices, provide insight into the link between weather/road conditions and maintenance decisions/outcomes. These models can be used in various ways—to understand how particular weather triggers specific actions (descriptive), to predict probable resource use given forecast conditions (predictive) or to provide guidelines as to optimal response to various conditions (normative). Ultimately, one would expect that descriptive models would be the basis for, or at least inform, the development of both predictive and normative models and indices. The current paper reports on a new approach for developing descriptive models of salt usage, using classification techniques and highway-segment specific data on an hourly basis.

The overall conclusion of the study is that classification trees show considerable promise in identifying the different weather conditions that result in specific salt application rates. This conclusion is based on the fact that the trees developed for the Highways 417 and 138 in the Vankleek Hill patrol near Ottawa identified both variables and threshold values that are consistent with expert knowledge (e.g., the need to apply salt prior to freezing rain) and scientific principles (e.g., ground frosts occur when the air and dewpoint temperature converge). Key to this success was the use of fine-scale salt usage data (the AVL data were easy to use and appeared accurate) and a wide range of weather and road condition variables from multiple data sources.

As well, it was determined that various types of classification trees can be used to predict the most likely hourly salt application rate for individual trucks or highway segments. Furthermore, these predictions can be used to estimate salt use at the daily or weekly level with a relatively high level of accuracy; when compared to conventional regression modeling this new approach produced r-squared values that are considerably higher—approximately 8 percent higher at the daily level (R^2 values of 0.69). Preliminary analysis at the weekly provides even better results.

In terms of objective three, the current study was not able to develop any empirical estimates of salt savings due to direct liquid application (DLA) of chemicals because of the nature of the data set used. In the three seasons analyzed, there were only 17 days on which DLA was applied and it was applied to only one of the seven highway segments examined. Thus the data set was too small to detect any effects. That said, it is important to note that the modeling structure used could accommodate DLA and other maintenance practices as additional explanatory variables and the expected result would be that DLA would lead to fewer treatment hours or possibly lower salt application rates than would otherwise have occurred.

RECOMMENDATIONS

Given the promising nature of the results, there are three recommendations for further work.

1. The data set used to develop the tree-based classifications should be expanded to include more patrols and/or seasons. Based on this expanded data set, a set of pruned trees that represent the most important conditions leading to salt applications of different amounts should be developed—initially for each individual highway segment, and then for all segments combined.

2. The results of the exercise should be shared with Ministry personnel and contractors who

are in charge of operational decisions, in order to explore the thinking behind certain types of decisions and to identify other explanatory variables that have not been included in the analysis and/or captured by the models.

3. Other explanatory variables that pertain to weather (e.g., the difference between the ambient and dew point temperature) and local factors (e.g., terrain, orientation, snow fencing) should be included in the analysis, where possible.

4. The list of explanatory variables should be expanded to include information on plowing, pre-wetting and DLA, as a way of exploring the combined or offsetting effects of different maintenance operations.

5. Residual analyses should be conducted throughout the modeling exercise in order to better understand those conditions when the models perform well and when the models do not perform well.

6. Different levels of aggregation should be explored so that model fit can be reported at different spatial-temporal units of analysis in ways consistent with different Ministry needs.

References

Adams, M.J., “Winter Severity Index”, Wisconsin DOT, internal document. 2001.

Andrey, J., Li, J. and Mills, B., “A winter index for benchmarking winter road maintenance operations on Ontario’s highways”, Proceedings of the Transportation Research Board’s 80th Annual Meeting, January 7-11, 2001, Washington, D.C.

Andrey, J., Brenning, A., Kirchhoff, D., Mills, B. and Perchanok, M., “Modeling winter maintenance activities using classification trees”, Transportation Research Board 4th Surface Transportation Weather and 7th Snow and Ice Conferences, June 16-19, 2008, Indianapolis.

Boselly, S., Thornes, J., Ulberg, C. and Ernest, D., “Road Weather Information Systems”, Unpublished research report, 1996.

Breiman, L., Friedman, H. J., R. A. Olshen, A. R. and Stone, J. C., *Classification and regression trees*. Wadsworth, California, 1984.

Breiman, L. 1996., “Bagging predictors”, Machine Learning, Vol. 24, No. 2, 1996, pp. 123-140.

Brenning, A., 2005. ‘Spatial prediction models for landslide hazards: review, comparison and evaluation”, Natural Hazards and Earth System Sciences, Vol. 5, 2005, pp. 853-862.

Brenning, A., Kaden, K. and Itzerott, S. ‘Comparing classifiers for crop identification based on multitemporal Landsat TM/ETM data”, In: M. Braun (ed.), Proceedings, Second Workshop of the EARSeL Special Interest Group on Remote Sensing of Land Use and Land Cover. 28-30, 2006, Bonn, Germany, pp. 64-71.

Brenning, A. and Lausen, B., “Estimating error rates in the classification of paired organs”, Statistics in Medicine, in press, 2008, DOI: 10.1002/sim.3310.

Bühlmann, P., “Bootstraps for time series”, Statistical Science, Vol. 17, 2002, pp. 52-72.

Carmichael, C., Gallus, W., Temeyer, B. and Bryden, M. “A winter weather index for estimating winter roadway maintenance costs in the Midwest”, Journal of Applied Meteorology, Vol. 43, 2004, pp.1783-1790.

Cornford, D. and Thornes, J., “A comparison between special winter indices and expenditure on winter road maintenance in Scotland” International Journal of Climatology, Vol. 18, 1996, pp.399-357.

Davison, A. C., Hinkley, V. D. and Young, A. G., “Recent developments in bootstrap methodology”.

Statistical Science, Vol. 18, 2003, pp. 141-157.

Decker, R., Bignell, J., Lambertsen, C. and Porter, K., “Measuring efficiency of winter maintenance practices”, Transportation Research Record 1741, 2001.

Efron, B. and Tibshirani, R., *An introduction to the bootstrap*, Chapman and Hall, New York, 1993.

Gustavsson, T., “Use of indices for classification of winter climate”, Proceedings of the 8th International Road Weather Conference, University of Birmingham, 1996.

Hand, D. J., *Construction and assessment of classification rules*. Wiley, Chichester, 1997.

Hand, D.J., Li, G. H. and Adams, M. N., “Supervised classification with structured class definitions”, Computational Statistics and Data Analysis Vol. 36, 2001, pp. 209-225.

Hastie, T.J. and Tibshirani, R., *Generalized Additive Models*. London: Chapman and Hall, 1990.

Hastie, T. J., Buja, A. and Tibshirani, R., “Penalized discriminant analysis”. Annals of Statistics, Vol. 23, No. 1, 1995, pp. 73-102.

Hothorn, T. and Lausen, B., “Bundling classifiers by bagging trees”, Computational Statistics and Data Analysis, Vol. 49, 2005, pp.1068-1078.

Hulme, M., “A new winter index and geographical variations in winter weather”, Journal of Meteorology 7, 1982, pp. 294-300.

Ihaka, R. and Gentleman, R., “R: A language for data analysis and graphics”, Journal of Computational and Graphical Statistics, Vol. 5, 1996, pp. 299-314.

Johns, D., “The MOORI: a new winter index for WRM”, Proceedings of the 8th International Road Weather Conference, University of Birmingham, 1996.

Jones, B., “Road maintenance costs in Canada in winter”, In Andrey, J. and Knapper, K. C. (eds). Weather and Transportation in Canada, Department of Geography Publication Series No. 55, University of Waterloo, Waterloo, ON, 2003, pp. 121-142.

Knudsen, F., “A winter index based on measured and observed road weather parameters” Proceedings of the 7th International Road Weather Conference, Seefeld, Austria, 1994.

Mahle, A. and Rogstad, G., “NORIKS: A winter index for Norwegian conditions”, Proceedings of the 11th International Road Weather Conference, Sopporo, Japna, 2002.

Mahoney, W.P.III, Bernstein, B. Wolff, J., Linden, J., Myers, W.L., Hallowell, R.G., Cowie, J., Stern, A.D., Koenig, G., Phetteplace, G., Schultz, P., Pisano, P.A. and Burkheimer, D., “FHWA’s maintenance decision support system pilot", Transportation Research Record 1911, TRB, National Research Council, Washington, D.C., 2005, pp. 133-142.

McCullouch, B., Belter, D., Konieczny, T. and McClellan, T., "Indiana winter severity index", Proceedings of the 6th International Symposium on Snow Removal and Ice Control Technology, Spokane Washington, 2004.

Morin, D. and Perchanok, M., "Road salt use in Canada", In Andrey, J. and Knapper, k. C. (eds). Weather and Transportation in Canada, Department of Geography Publication Series No. 55, University of Waterloo, Waterloo, ON, 2003, pp. 143-159.

Moguerza, M. and Munoz, A., "Support vector machines with applications", Statistical Science, Vol. 21, 2006, pp.322-336.

Nixon, W. and Qui, L., "Developing a storm severity index", Transportation Research Record 1911, 2005, pp. 143-148.

Peters, A., Hothorn, T. and Lausen, B., "Generalised indirect classifiers", *Computational Statistics and Data Analysis*, Vol. 49, 2005, pp. 849-861.

Rissel, M. and Scott, D., "Staffing of maintenance crews during winter months", *Transportation Research Record* 1019, 1985.

R Development Core Team, *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2006.

Synectics Transportation Consultants, University of Waterloo, Environment Canada and IMOS Inc., "Development of winter severity models for Canadian winter road maintenance", Final report to Transportation Association of Canada, 2006.

Suggett, J. Hadayeghi, Al, Andrey, J., Mills, B. and Leach, G., "Development of a winter severity index for salt management in Canada", Presented 86th Annual Meeting Preprint CD-ROM, *Proceedings of the Transportation Research Board Annual Meeting*, Washington, D.C., 2007.

Tibshirani, R. and Hinton, G., 1998. 'Coaching variables for regression and classification", Statistics and Computing, Vol. 8, 1998, pp. 25-33.

Thornes, J.E., "Cost-effective snow and ice control for the 1990s", Transportation Research Record 1387, TRB, National Research Council, Washington, D.C., 1993, pp. 185-1990.
Venables, W.N. and Ripley, B.D., *Modern Applied Statistics, 4th edition*. New York: Springer, 2002.

Venalainen, A., "Estimation of road salt use based on winter air temperature", Meteorological Applications, Vol. 8, 2001, pp. 333-338.

Appendix

Independent Variable	Valid Values	Data from Environment Canada and AMEC/MTO Weather conditions reported at beginning of the hour
WxVisibility	0-25	Visibility distance (km)
WxWindDir	Sin (0-360)	Direction of wind (decimal degrees based on 16 cardinal points)
WxWindSpeed	0-100	Wind speed measured at 10m height (km/h)
WxTemp	-40 to 40	Dry bulb air temperature (˚C)
WxDewpoint	-40 to 30	Dewpoint temperature (˚C)
WxRH	10-100	Relative humidity (%)
WxSLP	95-107	Sea level air pressure (kPa)
WxCloud	0-10	Cloud amount where 0 clear and 10 is overcast(tenths)
WxR	0-3	Reported occurrence and intensity of rain. Light (1), Moderate (2), Heavy (3)
WxRW	0-3	Reported occurrence and intensity of rain showers. Light (1), Moderate (2), Heavy (3)
WxL	0-3	Reported occurrence and intensity of drizzle. Light (1), Moderate (2), Heavy (3)
WxZR	0-3	Reported occurrence and intensity of freezing rain. Light (1), Moderate (2), Heavy (3)
WxZL	0-3	Reported occurrence and intensity of freezing drizzle. Light (1), Moderate (2), Heavy (3)
WxS	0-3	Reported occurrence and intensity of snow. Light (1), Moderate (2), Heavy (3)
WxSG	0-3	Reported occurrence and intensity of snow grains. Light (1), Moderate (2), Heavy (3)
WxIP	0-3	Reported occurrence and intensity of ice pellets. Light (1), Moderate (2), Heavy (3)
WxSW	0-3	Reported occurrence and intensity of snow showers. Light (1), Moderate (2), Heavy (3)
WxSP	0-3	Reported occurrence and intensity of snow pellets. Light

Independent Variable	Valid Values	Data from Environment Canada and AMEC/MTO Weather conditions reported at beginning of the hour
		(1), Moderate (2), Heavy (3)
WxFog	0,1	Reported occurrence of fog
WxBS	0,1	Reported occurrence of blowing snow
Wx6hrPam	0-50	Accumulated liquid-equivalent precipitation over 6-hour period in which hour occurs (mm). Trace reported as 0.05.
WxDlyRainAm	0-100	Accumulated daily rainfall in which hour occurs (mm). Trace reported as 0.05.
WxDlySnowAm	0-100	Accumulated daily snowfall in which hour occurs (cm). Trace reported as 0.05.
WxDlyPrecipAm	0-100	Accumulated daily liquid-equivalent precipitation in which hour occurs (mm). Trace reported as 0.05.
WxSnowDepth	0-100	Depth of snow on the ground measured in the morning (cm)
WxWsWatch	0,1	Winter storm watch in effect
WxWsWarn	0,1	Winter storm warning in effect
WxSnowWarn	0,1	Snowfall warning in effect
WxZrWarn	0,1	Freezing rain warning in effect
WxFlashWarn	0,1	Flash freeze warning in effect
WxBlowSnowWarn	0,1	Blowing snow warning
WxWsAdv	0,1	Advisory of potential for winter storm conditions in effect
WxSnowAdv	0,1	Advisory of potential for heavy snowfall, snowfall just below warning criteria, or changeover from rain to snow
WxZlAdv	0,1	Advisory of potential for freezing drizzle or rain below warning criteria
WxBlowSnowAdv	0,1	Advisory of potential for blowing snow
WxRadObs	0,1	Valid or invalid observation during event (assumed during non-event periods)
WxRadPrecipAny	0,1	Identification of precipitation echoes on the image that are moving towards and may reach study region
WxRadPrecipStudy	0,1	Precipitation echoes observed in study region
WxRadNature	0-2	Morphology of precipitation echoes observed in study

Independent Variable	Valid Values	Data from Environment Canada and AMEC/MTO Weather conditions reported at beginning of the hour
		region (0 no echoes, 1 cellular/scattered, 2 solid/continuous)
WxRadCov	0-3	Precipitation echoes cover 0% (0), <25% (1), 25-50% (2), or >50% (3) of the study region
WxRadIntensityAvg	0-4	Most common intensity of precipitation echoes observed in the study region. Precipitation echoes not observed in study area portion of image (0), 9-20db or 0.1-.3cm/hour (1), 20-29db or 0.3-.75cm/hour (2), 29-35db or 0.75-1.5cm/hour (3), >=35db or >=1.5cm/hour (4)
WxRadIntensityMax	0-10	Maximum intensity of precipitation echoes observed in the study region. Precipitation echoes not observed in study area portion of image (0), 9-20db or 0.1-.3cm/h (1), 20-25db or 0.3-.5cm/h (2), 25-29db or 0.5-.75cm/h (3), 29-32db or 0.75-1.0cm/h (4), 32-35db or 1.0-1.5cm/h (5), 35-38db or 1.5-2.0cm/h (6), 38-42db or 2-3cm/h (7), 42-45db or 3-4cm/h (8), 45-47db or 4-5cm/h (9), >=47db or >=5cm/h (10).
WxFxTonightTomorrowTemp	-40 to 40	Air temperature (˚C) forecast based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowTempTrend	-1,0,1	Air temperature (˚C) trend (falling, nothing stated, rising) based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowSnow	0,1	Precipitation (flurries, snow showers or snow) based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowRain	0,1	Precipitation (drizzle, rain showers, or rain) based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowMixed	0,1	Precipitation (freezing drizzle, freezing rain, or ice pellets) based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowSnowAmLow	0-40	Low snowfall amount (cm) based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowSnowAmHigh	0-40	High snowfall amount (cm) based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h

Independent Variable	Valid Values	Data from Environment Canada and AMEC/MTO Weather conditions reported at beginning of the hour
		next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowPop	0-100	Probability of precipitation (%) based on 1530h forecast by Environment Canada.. Maximum if different probabilities issued for sub-periods (e.g., morning/afternoon) or forms of precipitation (e.g., rain/snow). *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxFxTonightTomorrowBS	0,1	Blowing snow conditions based on 1530h forecast by Environment Canada. *Tonight* applies from 1800h to 0600h next day. *Tomorrow* applies from 0600h-1800h next day.
WxRoadEr16Tair	-40 to 40	Air temperature (˚C) at MTO RWIS ER-16 recorded closest to the top of the hour
WxRoadEr16Tdew	-40 to 30	Dewpoint temperature (˚C) at MTO RWIS ER-16 recorded closest to the top of the hour
WxRoadEr16Psituation	1-15	Precipitation situation reported at MTO RWIS ER-16 recorded closest to the top of the hour. Other (1), Unknown (2), No Precip (3), Unident Slight (4), Unident Moderate (5), Unident Heavy (6), Snow Slight (7), Snow Moderate (8), Snow Heavy (9), Rain Slight (10), Rain Moderate (11), Rain Heavy (12), Freezing Rain Slight (13), Freezing Rain Moderate (14), Freezing Rain Heavy (15)
WxRoadEr16Tsurface	-40 to 70	Pavement surface temperature (˚C) at MTO RWIS ER-16 recorded closest to the top of the hour
WxRoadEr16SurfaceStatus	1-14	Status (condition) of road surface recorded closest to the top of the hour. Other (1), Error (2), Dry (3), Trace (4), Wet (5), Chemical Wet (6), Ice Warning (7), Ice Watch (8), Snow Warning (9), Snow Watch (10), Absorption (11), Dew (12), Frost (13), Absorption at Dew 7 (14)