Cold Regions Science and Technology 59 (2009) 227–233

Contents lists available at ScienceDirect

Cold Regions Science and Technology

journal homepage: www.elsevier.com/locate/coldregions

Computer assisted avalanche prediction using electronic weather sensor data

P. Cordy *, D.M. McClung, C.J. Hawkins, J. Tweedy, T. Weick

Department of Geography, University of British Columbia, Canada

ARTICLE INFO

Article history:
Received 15 October 2008
Accepted 27 July 2009

Keywords:
Avalanche
Prediction
Electronic weather sensor
Nearest neighbour
Cross-validation

ABSTRACT

A nearest neighbour model was used to predict avalanches in two highway corridors in British Columbia, Canada. The model accepts hourly electronic weather sensor data dynamically to produce automated predictions of the probability that avalanches will occur in the next 12 h. Output includes a list of nearest neighbours calculated by a Euclidian metric which provides information on patterns of avalanche activity in similar situations in the past. New variables are automatically generated from the hourly interval sensor measurements, including information about accumulated precipitation and maximum and minimum temperatures. A jackknife cross-validation routine generates fitness statistics by selecting test cases that are not temporally autocorrelated. The avalanche prediction system described here was applied operationally in Kootenay Pass, near Salmo, BC, and also at Bear Pass, near Stewart, BC, where accuracies of 76 and 72% were achieved respectively.

1. Introduction

The goal of avalanche forecasting is to minimize three sources of uncertainty about the state of instability of the snow cover: spatial and temporal variability, incremental changes from snow and weather conditions, and the variability of human perception (McClung, 2002).

Updating the perception of instability and the associated risk to road traffic by human avalanche forecasters is a recursive process that continues through every day and night of the snow season. Though computer assisted avalanche prediction programs only partially represent reality, they can assist in providing access to relevant historical information and avalanche probability estimates that are as free as possible from human bias (McClung, 2000). The nearest neighbour algorithm (Buser, 1983; Purves et al., 2003; McClung and Tweedy, 1994) does this by providing the avalanche forecaster with weather and avalanche information for several days in the historic record that are similar in terms of weather. The nearest neighbour based Avalanche Forecast System (AFS) presented in this paper is more than 70% accurate despite data constraints imposed by the exclusive use of electronic meteorological station data. Previous efforts were based on manual observations (Floyer and McClung, 2003; Purves et al., 2003; McClung and Tweedy, 1994), though Roeger et al. (2003) showed that avalanche forecasts up to 24 h in advance could be usefully achieved by incorporating numerical weather prediction inputs. These advanced forecasts used a parameterized snow density proxy variable that was weakly correlated to manually recorded new snow density. The AFS does not require proxy variables. It only uses electronic data that are easily simulated by numerical weather forecasts, and therefore this simplifies the future application of numerical weather prediction to statistical avalanche forecasting.

It is critical that electronic data are incorporated into statistical avalanche prediction since regular manual observations are no longer taken in some areas, and automated weather stations now provide information for areas that are inaccessible during storms. This paper shows that electronic data alone can be sufficient to produce valid statistical avalanche predictions, and provides a new method for verification of hourly predictions using cross validation.

2. Description of locations

The AFS was applied to two widely separated and climatically distinct transportation corridors: Kootenay Pass on Highway #3 in southeast British Columbia, and Bear Pass on Highway #37A near Stewart, BC (Fig. 1, Table 1).

Kootenay Pass was the site of the first operational numerical avalanche prediction model in Canada (McClung and Tweedy, 1994), and is located on the Crowsnest Highway between Salmo and Creston in the Selkirk Mountains. Avalanches in Kootenay Pass are frequently initiated by explosives in order to clear out snow deposits and unstable layers, thus reducing the size of avalanche deposits on the road and shortening road closures. The timing, accuracy, and effectiveness of artificial avalanche triggering have evolved over time as the avalaunchers (Brennan, 2006) were replaced by fixed Gaz-Ex (Gubler, 1996) exploders. The Gaz-Ex can be fired remotely in any weather, at any time, during dry and moist avalanche periods. By contrast, most avalanches in Bear Pass occur naturally.

At Bear Pass, near the northwest coast of British Columbia, limited avalanche control is possible in inclement weather for some frequently

* Corresponding author. Tel./fax: +1 604 815 3789.
E-mail addresses: paulcordy@gmail.com (P. Cordy), mcclung@geog.ubc.ca (D.M. McClung).

0165-232X/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.coldregions.2009.07.006

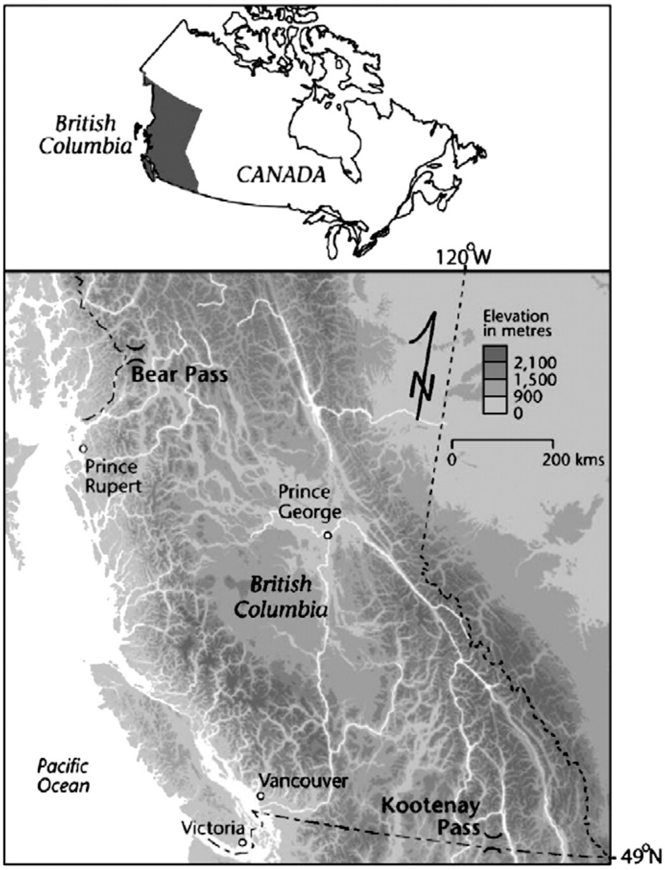

Fig. 1. Map showing the locations of study sites.

problematic paths by using a Howitzer cannon located near the summit of the pass. Otherwise, helicopters are the only way to deliver explosives to the start zones. Therefore, the road must be closed more often and for longer periods during persistent storms when helicopters cannot be used. This allows time for very large snow deposits to accumulate. This increases the severity of avalanche activity and prolongs efforts at clearing avalanche debris from the road.

Table 1 summarizes some important climatological, geographic, and operational features of Kootenay and Bear Passes. For more thorough descriptions of these sites refer to McClung and Tweedy (1994), Floyer and McClung (2003), and Roeger et al. (2003).

3. Data

Data relevant to avalanche prediction can be divided into three classes according to LaChapelle (1980) and McClung and Schaerer (2006): stability (class I), snowpack (class II), and meteorological (class III). These classes, described in Table 2, are numbered in order of decreasing informational entropy, which is defined as the relevance and ease of interpretation of a given piece of evidence.

Electronic precipitation, snow depth data, and air temperature at both locations are measured using a liquid precipitation gauge, a sonic ranging snow height sensor, and a digital thermometer placed on a tripod in the weather plot at road elevation. Wind data are taken by anemometers on nearby ridges. The data are collected and reported hourly through the British Columbia Ministry of Transportation's (MoT) Snow Avalanche and Weather System (SAWS) database. Data were filtered to exclude observations with missing data, periods in which one or more sensors were not properly functioning, and periods for which avalanche observations are not available electronically.

4. Methods

4.1. Nearest neighbours

When generating a prediction, the nearest neighbour algorithm chooses k rows of weather data from the historical database that are geometrically nearest to the current or forecast weather. Distance is computed in an n-dimensional predictor space where each dimension corresponds to a different input variable (such as new snow height, air temperature, total snow depth, wind speed). The proportion of positive neighbours (hereafter: NN ratio) is the number of neighbours that are associated with avalanches, divided by k. If the NN ratio exceeds the chosen threshold value (hereafter: *threshold k*), then the model predicts that avalanches will occur.

The NN ratio represents a relative estimate of the probability of avalanches given recent weather. These probability estimates are relative in the sense that larger numbers indicate a greater chance that avalanches will occur, but the optimal warning level (*threshold k*) can vary depending on the forecast region. McClung and Tweedy (1994) found that the Kootenay Pass nearest neighbour model *threshold k* of 0.2 corresponds to a parametric linear discriminant analysis prediction warning level of 0.6, but in this study the NN ratio output was not rescaled using this equivalency. Instead, the AFS reports the raw NN ratio and lets the human forecaster interpret the meaning of this value relative to the *threshold k*.

Nearest neighbours distances in this study are computed in the Euclidian distance metric (Buser, 1983). As in McClung and Tweedy (1994), and Floyer and McClung (2003), this study contains a logarithmic transform $X_i = \ln(X_i + 1)$ (Bovis, 1976) to remove positive skewness from the precipitation data. All variables are then standardized by subtracting the overall mean and dividing the result by the standard deviation.

Values of k and *threshold k* are taken from McClung and Tweedy (1994) and Floyer and McClung (2003); $k = 30$ in both locations, with thresholds of 6 and 7 respectively.

When the nearest neighbour method is applied to the hourly data, the k nearest neighbours often include several observations from the same day. This effectively reduces the number of unique weather events upon which the prediction is based, and gives greater weight to days from which several neighbours are drawn. Therefore only the most similar (in terms of Euclidian distance) neighbour to be chosen from any calendar day is used in each prediction, and all other observations from that day are ignored.

Table 1
Climatic and operational characteristics of Kootenay and Bear Passes (Haegli, 2004; Floyer and McClung, 2003; McClung and Schaerer, 2006).

	Kootenay Pass	Bear Pass
Snow climate	Transitional	Maritime
Average annual snowfall	908 cm/year	805 cm/year
Range of differences in elevation between the highway and start zones	100 m–600 m	100 m–1900 m
Roadside weather plot elevation	1774 m	370 m
Avalanche control methods/vehicles currently used	Gaz-Ex, helicopter, hand charges	105 mm recoilless rifle, helicopter, hand charges
Length of electronic data record	1997–2007	2001–2007

Table 2
Classes of informational entropy and their characteristics (high entropy data are more locally relevant and interpretable).

Class	Examples	Characteristics	Rationale for inclusion in or omission from the AFS
Stability (I)	Propagating cracks Recent avalanche activity	o Locally relevant o Easy to interpret	o Limited spatial support o Not detected electronically
Snowpack (II)	Snow layer density profile Weak layer depth Weak layer grain type Temperature gradient	o Subsurface information o Highly spatially variable o Recorded in snow profiles	o Strongly affected by terrain, elevation, and aspect o Not detected electronically o Snow profiles are taken infrequently at different locations
	New snow density Surface penetration	o Measured using penetrometer or foot penetration, or mass/volume ratio o Computed by proxy	o Density is not directly detected electronically o Electronic sonic ranging snow depth sensor underestimates low density snow, making computed densities inaccurate
Meteorological (III)	Wind speed Temperature New snowfall Water equivalent of new snow	o Relevant at synoptic scale o Measured electronically and manually	o Hourly measurements enable recent updates of avalanche predictions on demand o Easy to automate

Data used in this study are shaded.

4.2. Automated data manipulation

The AFS retrieves current weather sensor data hourly and the nearest neighbour algorithm makes predictions by comparing the new information to the historical data. These historical datasets are generated using raw hourly historical data from SAWS. Each row of current and historical data contains both raw and memory variables. Raw variables use the original values provided by SAWS, such as present temperature and precipitation during the last hour. Memory variables, which retain information about accumulation, differences, or extreme values within a given period of observation, are computed as in Table 3. The observation period defines the number of hours previous to the current observation time over which memory variables are computed. Furthermore, each hourly datum is labeled "true" if avalanches of any moisture type (dry, moist or wet) occurred within 12 h of that datum and "false" otherwise (Cordy et al., 2006). In this paper we refer to this 12 h interval as the response period.

The new snowfall variable is computed by finding the change in snow height over the observation period. This method implicitly includes snow settlement into the model, and also reduces the error caused by the snow sensor's imprecision in low density snow. The liquid precipitation gauge data are used to measure the water equivalent of new snow by summing hourly values over the same observation period as the new snowfall variable.

Table 4 shows the variables used in this study, as well as the observation periods over which each memory variable is computed. These observation periods were optimized by trial and error with reference to the fitness statistics listed in Section 4.4. For each memory variable, observation period values from 1 to 48 hours were tested in one hour intervals, and those that produced the highest performance scores were used in the final model.

4.3. Jackknife cross-validation

In the nearest neighbour context, predictions are made for each "test" datum by drawing neighbours from a separate "training" dataset. McClung and Tweedy (1994) and Floyer and McClung (2003) reserve 20% of historic data rows for model verification and the remaining data are used to find the parameters of the discriminant model.

Cross-validation (Hastie et al., 2001) is an iterative process that uses the entire dataset for fitness testing while still maintaining separate "training" and "testing" datasets in each iteration. A simple cross-validation extension of this process would divide the data into five blocks, each containing 20% of the data. In the first iteration, the first block of data is used as test data and the remaining four blocks are training data. In the next iteration, the second block of data is used as test data, and the training data come from the first, third, fourth, and fifth blocks. This process is repeated until all blocks have been tested.

Jackknife cross-validation only uses one test observation in each iteration, instead of a block of data. The rows adjacent to the test observation are excluded from the training data, and the number of excluded rows amounts to ≥20% of the dataset. Since each observation row is associated with the 12 h of avalanche occurrences that follow, testing hourly data in chronological sequence effectively tests predictions of the same avalanches multiple times. Therefore, every 17th row is tested. This "cross-validation sampling period" value of 17 ensures that the periodicity of rows chosen for testing is greater than the 12 h response period and does not match diurnal or monthly cycles.

If p is the cross-validation sampling period, m is the test data exclusion mask, and t_i is the current observation row, (indexed by i), then Algorithm 1 defines the cross-validation scheme used in this study.

Algorithm 1
Jackknife cross-validation.

1. For row t_i, make a training dataset that excludes all data within +/- m hours of time t_i.
2. Make a prediction for row t_i.
3. Iterate for $t_{i+1} = t_i + p$.

Table 3
AFS memory variable generation functions for a 24 h observation period.

Function	Value returned	Sample sensor	Sample variable
Max/min	Maximum or minimum value that occurs during observation period.	Thermistor	24 h maximum temperature
Sum	Sum of all hourly values recorded during the observation period	Liquid precipitation gauge	24 h new precipitation
Difference	The net change in value of the variable over the observation period	Sonic ranging snow height sensor	24 h new snowfall

Table 4
Variables used in the analysis at both field sites, and the sensors that generate each variable (N/A indicates that no such memory variable is used in the analysis).

Sensor		Kootenay Pass variables	Bear Pass variables
Sonic ranging	Raw	Total snow depth	Total snow depth
snow height	Memory	N/A	30 h new snowfall
Precipitation	Raw	Hourly precipitation	Hourly precipitation
gauge	Memory	24 h new precipitation	30 h new precipitation
Anemometer	Raw	Wind speed	N/A
	Memory	N/A	18 h maximum wind speed
Thermometer	Raw	Present temperature	Present temperature
	Memory	24 h maximum temperature	N/A

Table 5
Contingency table (total number of cases $N = a + b + c + d$).

		Observation	
		Avalanche	No Avalanche
Prediction	Avalanche	a: true positive	b: false positive
	No Avalanche	c: false negative	d: true negative

4.4. Fitness metrics

Fitness metrics are used to evaluate and compare the accuracy and performance of models. There are four fitness metrics used in this study: proportion correct (PC), unweighted average accuracy (UAA), bias, and the Peirce skill score (PSS). The equations that define these fitness metrics are shown in Tables 5 and 6 (Jolliffe and Stephenson, 2003).

The proportion correct (PC) is an intuitively appealing measure of the overall success of the model; it is the proportion of the whole sample that is a correct forecast of either event. It is included here because it is commonly used in numerical avalanche forecasting research. Jolliffe and Stephenson (2003) argue that it is an unreliable measure of performance in part because when predicting rare events like avalanches, the best PC is often achieved by always forecasting non-events (0.87 and 0.83 for Kootenay and Bear Passes respectively, see Table 7). The unweighted average accuracy (UAA) is the most informative accuracy measure to use for rare events such as avalanches because it gives equal weight to events and non-events (Purves et al., 2003).

Peirce's skill score, also known as the Hanssen–Kuipers discriminant in Purves et al. (2003) and true skill score (Roeger et al., 2003), measures skill relative to an unbiased random reference forecast, and therefore is the most important performance metric in this study. A score of 0 indicates that the model has the same forecast skill as a random forecast or a constant prediction of event or non-event. A perfect forecast would have a PSS of 1 (Jolliffe and Stephenson, 2003).

Bias values greater than (less than) one indicate that the event was forecast more than (less than) it was observed. This is known as overforecasting (underforecasting) (Roeger et al., 2003).

5. Results

Table 7 contains the fitness statistics associated with the previous models and the AFS at both Bear and Kootenay Passes. Note that the AFS and the previous Bear Pass model (Floyer and McClung, 2003) make predictions of all avalanche types together (moist, wet and dry), whereas McClung and Tweedy (1994) discriminated between dry and moist/wet types and applied a different model and variable set to each.

These results are similar to those obtained by McClung and Tweedy (1994) and Floyer and McClung (2003) using twice daily manual observations. PSS values indicate that all of the models listed are more skillful than a random forecast or a constant forecast of non-avalanche occurrence. Values for PC and UAA for any given model (excluding Purves et al., 2003) are nearly equal despite the fact that non-avalanche days, by virtue of their greater number, contribute more to the PC value.

Table 6
Fitness metrics, ranges, and perfect score values.

Fitness metric	Equation	Range	Perfect score
Proportion correct (PC)	$\frac{a+d}{N}$	0–1	1
Unweighted average accuracy (UAA)	$0.5\left(\frac{a}{a+c}+\frac{d}{b+d}\right)$	0–1	1
Peirce's skill score (PSS)	$\frac{a}{a+c}-\frac{b}{b+d}$	−1 to 1	1
Bias	$\frac{a+b}{a+c}$	0–∞	1

Table 7
Fitness statistics for various models.

Area	PC	UAA	PSS	Bias	Number of test rows	# of test rows associated with avalanches
Kootenay Pass (AFS)[b]	0.75	0.76	0.54	2.12	1336	170
Kootenay Pass (McClung NN)[b]	0.72	0.71	0.41	1.97	784	92
Kootenay Pass (McClung LDA)[b]	0.77	0.75	0.52	1.96	784	92
Bear Pass (AFS)[b]	0.72	0.72	0.46	2.30	1380	230
Bear Pass (Floyer)[b]	0.72	0.72	–	–	–	–
Scotland (Purves)[b]	0.80	0.76	0.52	1.12	202	51

[b] AFS refers to the model presented here, Floyer refers to Floyer and McClung (2003), McClung refers to McClung and Tweedy (1994), NN refers to the nearest neighbour model, LDA refers to the linear discriminant analysis model, and Purves refers to Purves et al. (2003).

This suggests that the model predicts avalanche days and non-avalanche days with equal accuracy despite the fact that the bias indicates that all models over forecast avalanches.

A partial explanation of the high bias can be seen by considering the time series of predictions and avalanche events shown in Fig. 2. In this graph, the intensity and duration of avalanche periods are indicated by the avalanche activity index (AAI), which is defined as the sum of all of the sizes of avalanches (according to the Canadian size classification system, CAA, 2002) in a given period (McClung and Tweedy, 1994).

For Kootenay Pass, the best model fitness is achieved when a threshold of 6 out of 30 neighbours is used (NN ratio = 0.2). Explosive avalanche control measures (using Gaz-Ex or bombs thrown from a helicopter) impact the performance of the model, as illustrated by the shaded bands in Fig. 2. As an example, in the shaded period from January 16 to 20, the NN ratio indicated by the model remains high despite the fact that snow deposits in the start zones were cleared out using explosives at the beginning of the storm. As a result, there is no avalanche activity in the rest of the storm cycle and therefore the AFS counts each NN ratio above the threshold of 0.2 as a false alarm even though there could perhaps have been avalanches during that period if there had not been explosive control of deposits. This phenomenon was observed repeatedly over the course of the winter of 2006/7 when the AFS was tested in the field at Kootenay Pass. During verification, this phenomenon contributes false alarm predictions that degrade the measured fitness of the models.

In Fig. 2, as in most such time series generated by the AFS, it appears that the high values of NN ratio (>0.5) given by the AFS are almost always associated with avalanche occurrences. The majority of false alarms that are not attributable to recent explosive avalanche control occur when the NN ratio is between 0.2 and 0.4 (such as in periods highlighted by dashed boxes). Given that a datum is labeled an avalanche period if the NN ratio is greater than 0.2, it is clear that false alarms where the NN ratio is between 0.2 and 0.4 contribute significantly to the high bias of the AFS models (Fig. 2). However, since lower (higher) values of NN ratio seem to indicate higher degrees of uncertainty (certainty) that avalanches will occur, the NN ratio can justifiably be used as an estimate of the relative probability of avalanche occurrences.

Also, any given range of NN ratios (for example, between 0.3 and 0.4) is usually within 0.1 of the posterior probability of avalanches for that range of predictions. Fig. 3 shows that predicted NN ratios are lower than the posterior probability of avalanches for a given NN ratio, but most points plot fairly close to the perfect forecast line. It is interesting to note that for the Kootenay Pass model, where explosive triggered avalanches represent a far larger proportion of total avalanche occurrences than in Bear Pass, NN ratios near 0.5 and 0.6 lie farther from the perfect forecast line. This further supports the argument that fitness scores are negatively impacted by the increased

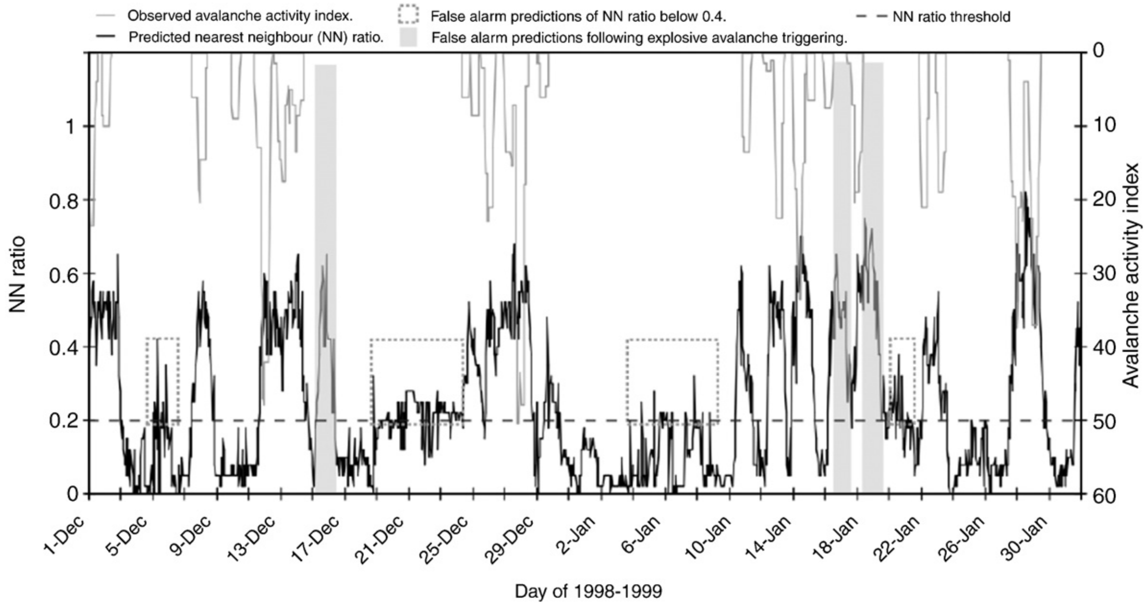

Fig. 2. Time series from Kootenay Pass, showing the timing of observed avalanches and the nearest neighbour ratio. The avalanche activity index refers to the sum of avalanche sizes in the response period, and the dashed line marks the warning level, or threshold k above which the period is classified as an avalanche period. The dashed boxes highlight periods of false alarms are due to low nearest neighbour ratio values which still exceed the threshold, and the gray bands highlight false alarms due to pre-emptive avalanche triggering.

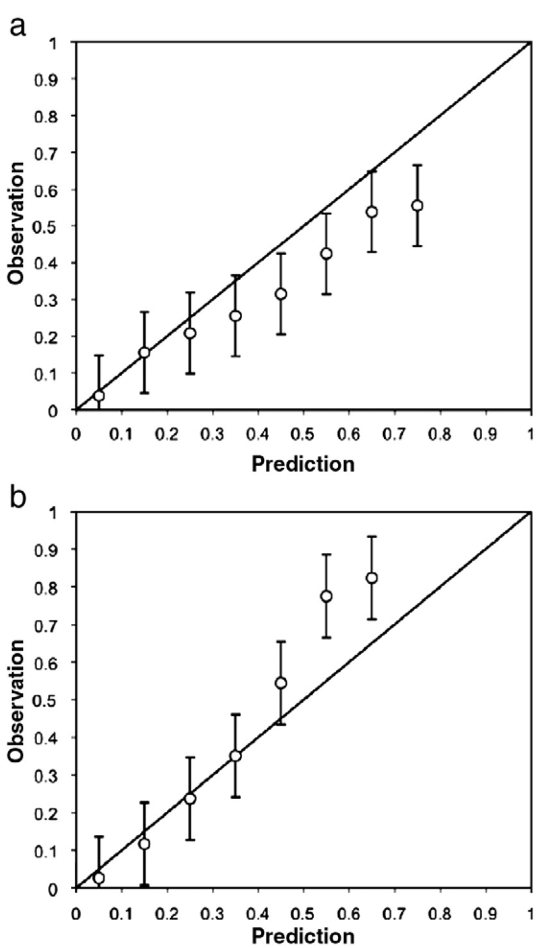

Fig. 3. Comparison of predicted nearest neighbour ratios with posterior probability of avalanches for Kootenay Pass (a) and Bear Pass (b). Prediction events are grouped into ranges of predicted nearest neighbour ratios (0 to 0.1, >0.1 to 0.2, etc.) and the posterior probabilities are plotted in the middle of each range. The diagonal indicates perfect skill. Two standard deviations from perfect skill are shown as error bars.

incidence of false alarms when avalanches are triggered with explosives early in the storm cycle.

Fig. 4 shows an entire season of AFS predictions and observed avalanche occurrences at Kootenay Pass. At this resolution it is clear that the AFS can recognize major avalanche periods. In many seasons tested, despite the fact that the datasets from which neighbours are drawn include moist and wet avalanches as well as dry avalanches, there are commonly a significant number of avalanches in spring that are not forecasted. The AFS is unable to predict moist and wet avalanche events accurately, even when using similar variables to those used in McClung and Tweedy's (1994) wet avalanche specific model.

6. Discussion

The Avalanche Forecast System described in this paper is integrated with the BC Ministry of Transportation electronic sensor data systems, and is sufficiently flexible to be tested province-wide. Due to constraints on time and programming resources, the variables used in this analysis were not weighted. Purves et al. (2003) computed variable weights that reflect the relative importance of predictors using a genetic algorithm, and this improved accuracy by 5% over an un-weighted model with normalized data. In order to account for correlations among variables, the Mahalanobis distance metric for computing nearest neighbours was used previously with BC Ministry of Transportation data (McClung and Tweedy, 1994; Floyer and McClung, 2003). Therefore, accuracy might be improved at both sites by using a different distance metric, or by adding weights to the predictor variables (Friedman, 1994) and optimizing them stochastically (Purves et al., 2003).

Previous models also included new snow density or penetration information. These variables proved to be statistically significant in previous studies (McClung and Tweedy, 1994; Floyer and McClung, 2003). Roeger et al. (2003) used a temperature-based proxy function to replace density variables used in the McClung and Tweedy (1994) model, but this function did not improve accuracy in the model presented here.

The Avalanche Forecast System created for the MoT has potential to be a useful addition to the suite of tools provided to the avalanche technicians. During operations, technicians at Kootenay Pass refer to the predictions of the AFS to support their inductive reasoning regarding the state of instability of the snow cover (K. J. Malone,

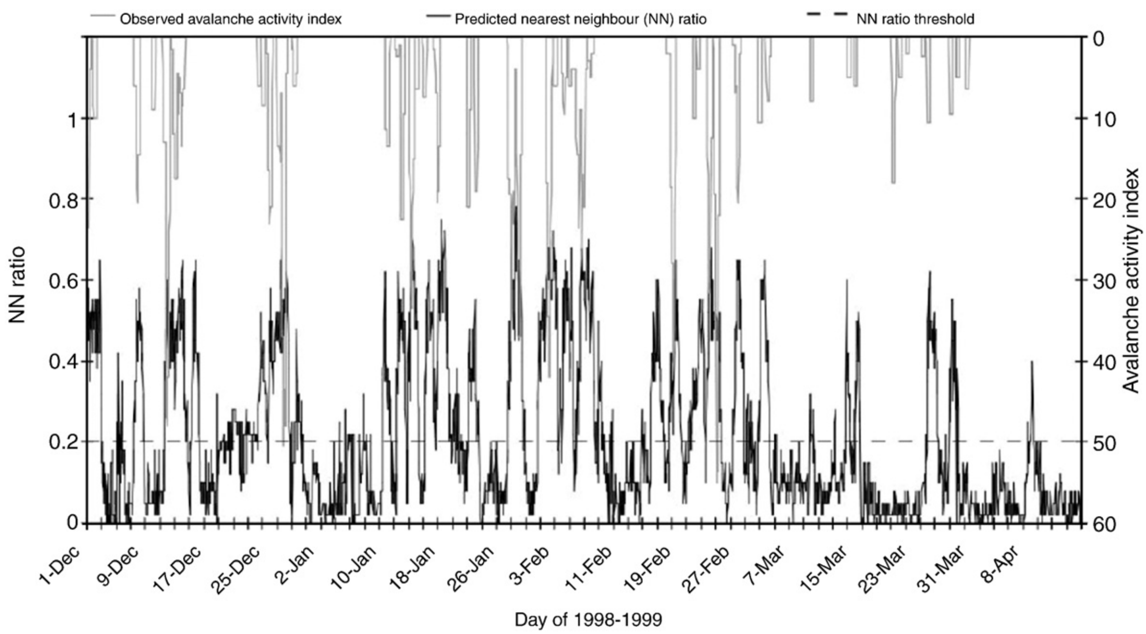

Fig. 4. Time series from Kootenay Pass showing the timing of observed avalanches and the nearest neighbour ratio. The avalanche activity index refers to the sum of avalanche sizes in the response period, and the dashed line marks the warning level, or threshold *k*, at which the period is classified as an avalanche period.

personal communication 2007). They use the nearest neighbour distances to judge the degree of uncertainty of the numerical predictions and they review the avalanche occurrences listed in the nearest neighbour report to confirm and improve their intuition regarding patterns of weather and avalanche occurrences.

The forecast resolution is 12 h since each nearest neighbour is associated with avalanches that occurred within 12 h after the time of that neighbour, and some variables are cumulative measures over longer periods. However, the hourly output resolution enables forecasters to view predictions based on the most up to date information available, at any time of day. Hourly output resolution could also enable the system to track changes in the likelihood of avalanches that result directly from weather changes that take place in fewer than 12 h, but further study is required to assess the prediction skill of hourly changes in AFS output.

Future AFS developments will deploy the software in other BC highways avalanche mitigation programs, and extend forecasts farther into the future using numerical weather prediction output as in Roeger et al. (2003).

7. Conclusions

The new Avalanche Forecasting System presented here predicts avalanches in the coming 12 hours using hourly interval electronic weather sensor inputs to a nearest neighbour model. The system achieved unweighted average accuracies of 72% for Bear Pass and 76% for Kootenay Pass. Despite being built without the use of new snow density and penetration information, these models still compare favorably with previous studies based on twice daily manual observations.

Model fitness is degraded slightly in areas where avalanches are more frequently triggered with explosives, but both models presented here show good skill. This shows that there is promise for the use of electronic weather sensor information for avalanche prediction in different climatic regions across British Columbia. This could be especially useful for MoT avalanche technicians in training, and for experienced technicians who are responsible for numerous and widely separated avalanche areas. It remains to be seen if variation in AFS predictions at hourly scale capture hourly scale changes in avalanche probability, however the system is now amenable to immediate incorporation of new snowpack or stability information as it is recorded, and predictions are available to human forecasters at their convenience. Furthermore, numerical weather prediction output can replace all electronic weather sensor variables used in this study, allowing extension of avalanche predictions farther into the future without parameterization of missing snow density information.

The jackknife cross-validation function created for this study is vital for generating defensible accuracy and performance values when considering hourly predictions. This function ensures that all test cases are independent of each other and that 20% of the data surrounding the test case are excluded from the set of potential neighbours.

Acknowledgments

This work is supported by the Canadian Mountain Holidays; the Natural Sciences, Engineering and Research Council; and the British Columbia Ministry of Transportation. This research was made possible by the guidance and assistance of Dr. Roland Stull of UBC Earth and Ocean Science, and Kevin Maloney of the British Columbia Ministry of Transportation. Stagleap Provincial Park provided the essential field experiences that enriched this project.

References

Bovis, M.J., 1976. Statistical forecasting of snow avalanches, San Juan Mountains, southern Colorado. Journal of Glaciology 18 (78), 83–130.
Brennan, J., 2006. The evolution of the avalauncher. Proceedings of the International Snow Science Workshop 2006. Telluride, USA. pp. 572–575.
Buser, O., 1983. Avalanche forecasting with the method of nearest neighbours: an interactive approach. Cold Regions Science and Technology 8 (2), 155–163.
CAA, 2002. Observation Guidelines and Recording Standards for Weather, Snowpack and Avalanches. Canadian Avalanche Association, Revelstoke, Canada. pp. 78.
Cordy, P., McClung, D.M., Hawkins, C.J., Weick, T., 2006. Automated numerical prediction using electronic meteorological and manual snowpack data. Proceedings of the International Snow Science Workshop 2006. Telluride, USA. pp. 612–617.
Floyer, J.A., McClung, D.M., 2003. Numerical avalanche prediction: Bear Pass, British Columbia, Canada. Cold Regions Science and Technology 37 (3), 333–342.
Friedman, J.H., 1994. Flexible metric nearest neighbour classification. Stanford University, Department of Statistics: Technical Report. http://www-stat.stanford.edu/~jhf/ftp/flexNN.pdf.
Gubler, H., 1996. Remote avalanche warning-, alarm- and control systems, fundamentals, applications and experience. Proceedings of the International Snow Science Workshop, Banff. pp. 165–172.
Haegli, P., 2004. Scale analysis of avalanche activity on persistent snowpack weaknesses with respect to large scale avalanche forecasting. University of British Columbia Dept. of Earth and Ocean Sciences, PhD Thesis.

Hastie, T., Tibshirani, R., Friedman, J., 2001. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer-Verlag, New York.
Jolliffe, I., Stephenson, D.B. (Eds.), 2003. Forecast Verification: a Practitioner's Guide in Atmospheric Science. Wiley, Chinchester, England.
LaChapelle, E.R., 1980. The fundamental processes in conventional avalanche forecasting. Journal of Glaciology 26 (94), 75–84.
McClung, D.M., 2000. Predictions in avalanche forecasting. Annals of Glaciology 31, 377–381.
McClung, D.M., 2002a. The elements of applied avalanche forecasting, Part I: The human issues. Natural Hazards 26 (2), 111–129.
McClung, D.M., Schaerer, P., 2006. The Avalanche Handbook3rd edition. The Mountaineers, Seattle WA, U.S.A. 342 pp.
McClung, D.M., Tweedy, J., 1994. Numerical avalanche prediction: Kootenay Pass, British Columbia. Journal of Glaciology 40, 350–358.
Purves, R., Morrison, K.W., Moss, G., Wright, D.S.B., 2003. Nearest neighbours for avalanche forecasting in Scotland — development, verification and optimisation of a model. Cold Regions Science and Technology 37 (3), 343–355.
Roeger, C., McClung, D.M., Stull, R., 2003. Verified combination of numerical weather and avalanche prediction models at Kootenay Pass, British Coloumbia, Canada. Annals of Glaciology 38, 334–346.