

Modeling a System for Decision Support in Snow Avalanche Warning Using Balanced Random Forest and Weighted Random Forest

Sibylle Möhle^{1,2}, Michael Bründl², and Christoph Beierle¹

¹ Dept. of Computer Science, University of Hagen, 58084 Hagen, Germany
`smoehle@acm.org`, `beierle@fernuni-hagen.de`

² WSL Institute for Snow and Avalanche Research SLF, 7260 Davos Dorf, Switzerland
`bruendl@slf.ch`

Abstract. In alpine regions, traffic infrastructure may be endangered by snow avalanches. If not protected by physical structures, roads need to be temporarily closed in order to prevent fatal accidents. For assessing the danger of avalanches, local avalanche services use, amongst others, meteorological data measured on a daily basis as well as expert knowledge about avalanche activity. Based on this data, a system for decision support in avalanche warning has been developed. Feasible models were trained using Balanced Random Forests and Weighted Random Forests, yielding a performance useful for human experts. The results are discussed and options for further improvements are pointed out.

Keywords: Balanced Random Forest, Weighted Random Forest, avalanche warning, decision support system, rare events, class imbalance

1 Introduction

Snow avalanches pose a serious threat in alpine regions. They may cause significant damages and fatal accidents. Therefore, local avalanche services responsible for avalanche safety in communities and for traffic infrastructure have been established in alpine countries. Their task is to protect people from the impact of snow avalanches by temporary measures, like the closing of roads, ordering people to stay in buildings, evacuation, or artificial avalanche triggering [20]. Thus, assessing the local risk of snow avalanches is of vital importance, and requires expert knowledge, intuition, and process understanding. Decision support systems such as NXD2000 [9, 10] based on the method of nearest neighbors [4] help local avalanche forecasters to base their decisions in addition to their knowledge and experience on more objective criteria. Precipitation (new snow or rain), wind, air temperature and solar radiation are the main factors influencing the formation of avalanches. Local avalanche forecasters base their daily judgment of avalanche danger on a careful analysis of meteorological variables and snowpack properties influencing the stability of the snowpack. This assessment relies

heavily on a sound understanding of the physical processes in the snowpack but also on experience and comparisons with similar situations observed in the past. When using NXD2000, the ten days being most similar to the current situation and the avalanche activity that occurred within the corresponding time slots are selected by the program and presented to the user. The output of the model provides additional evidence that a similar avalanche activity might take place in a specific situation.

Classification and regression trees [3] were applied in [18] for forecasting large and infrequent snow avalanches. The work reported in [18] concentrated on one single avalanche path. In [12], classification trees were adopted for forecasting avalanches in coastal Alaska; the region accounted for comprised over 100 avalanche paths, and prior to modeling, variable selection was performed. In [17], wet-snow avalanches were predicted using classification trees and Random Forests [2]. In [12, 17, 18], the training data were selected using statistical methods.

The purpose of the work reported here is to develop a decision support system for assessing the local hazard of snow avalanches, based on the data collected within NXD2000. As local knowledge is essential and generalization to other locations is very difficult, the system has been developed for a specific area, the Canton of Glarus in Switzerland.

A second objective of our work was to investigate the suitability of Random Forests and variants thereof for modeling a decision support system for snow avalanche warning. Since Random Forests [2], which are used in [17], are suited for classification problems where the dependencies between the variables are unknown and non-linear, they are a candidate for modeling our system. However, in the given data set provided by NXD2000, avalanches represent rare events, and Random Forests are biased towards the majority class [5]. Due to the class imbalance, bootstrap samples drawn for decision tree construction may contain few or no examples from the minority class, hence the resulting decision tree will perform poorly on examples from the minority class [5]. Therefore, we employ two variants of Random Forests, Balanced Random Forest and Weighted Random Forest [5], and study their suitability for our application.

A third goal of our work is the elaboration of quality measures for the obtained models from the point of view of the avalanche service which is in charge of assessing the local avalanche danger. While we also use the quality measures employed in the work cited above, we propose to use positive and negative predictive values as additional quality measures for snow avalanche warning, providing an assessment of the forecast probability. It turned out that these measures are particularly useful for the human experts at the avalanche service of the Canton of Glarus.

The rest of this paper is organized as follows: In Section 2, we briefly address the problem of forecasting rare events. In Section 3, we define six different measures for assessing the quality of models and discuss their relevance for our application scenario. The weather and snow data being used and the variables derived from these data are presented in Section 4. The resulting models ob-

tained by employing Balanced Random Forest and Weighted Random Forest are described and discussed in Section 5 and Section 6. In Section 7 we conclude and point out further work.

2 Forecasting Rare Events

In the given data covering more than 40 years, avalanches represent rare events (i.e. 53 days with avalanches, but 6889 without an avalanche). A model predicting always the negative class, i.e. non-avalanches, achieves a high overall accuracy. However, it would be of no use as a decision support for avalanche warning, since the important cases are the positive ones, i.e. the avalanche days. Fatal accidents may be the consequence of roads not being closed due to a missed avalanche forecast. Therefore, with rare events, the overall accuracy of a classifier is not an adequate quality criterion [19].

Sampling as well as cost-sensitive learning are possible solutions to the problem of predicting rare events. Undersampling the negative class may result in the loss of important information, while oversampling the positive class may introduce duplicates of positive examples into the training data. This, in turn, bears the risk of learning specific examples [19].

When taking measures such as temporary road closures, resulting costs have to be considered. They consist of business interruption costs for the regional economy due to road closure as well as the efforts for road closure and opening and clearance of avalanche debris on the road. In the case of false negative forecasts, costs may be significant due to damages and fatal accidents. Costs need not only be monetary: With every false prediction, the avalanche service loses credibility. For this reason, it is not only important to achieve a low number of false negatives, but the number of false positives has to be low, too.

With cost-sensitive learning, different forecast types are assigned different costs. An ideal classifier minimizes the associated cost function [6]. Costs are case-specific and their estimation is difficult. Long-time damage statistics allow for a quantitative estimation of costs. Since appropriate data were not available in our application, a qualitative estimation of costs from a regional economic view was conducted. The following assumptions apply: With a positive forecast, the affected road section is closed. Considered costs include damage to persons and property and business interruption costs as well as loss of credibility in case of unnecessary closure; hence, costs for correct predictions were set to 0. According to the remarks made in the last paragraph, the costs for false negatives have to be higher than the costs for false positives. Cost-sensitivity can be achieved by assigning different weights to the positive and negative class [6].

In [5], variants of Random Forests [2] suited for the classification of rare events were proposed.

2.1 Balanced Random Forest (BRF)

Balanced Random Forest approaches the problem of class imbalance by constructing each decision tree from equally-sized bootstrap samples from the nega-

tive and the positive class. This ensures that positive and negative examples are both included in the training data set. The trees are grown using the CART algorithm [3] without pruning. The determination of the best split in each node is carried out analogously to Random Forests by testing a previously fixed number of randomly chosen variables.

2.2 Weighted Random Forest (WRF)

Weighted Random Forest implements cost-sensitive learning by assigning weights to classes. By assigning a higher weight to the positive class, misclassification costs for positive examples are higher and positive examples therefore gain weight in the training process.

3 Model Assessment

Forecasting an avalanche can be considered as a classification problem. The positive class contains the avalanche days, and in this work, it is assigned the value 1. The negative class contains the non-avalanche days and is assigned the value 0. The results can be represented in contingency tables as shown in Fig. 3.

		Observed	
		0	1
Predicted	0	TN	FN
	1	FP	TP

Fig. 1. Contingency table for event forecasting: TN denotes the number of true negative forecasts, FN the number of false negatives. The number of false positive and true positive forecasts are denoted by FP and TP , respectively.

The number of true negative forecasts is abbreviated as TN and denotes the number of cases in which neither an avalanche was predicted nor an avalanche occurred. Accordingly, the number of correctly predicted avalanches is abbreviated as TP . The number of false negatives refers to the number of missed avalanches and is abbreviated as FN , the number of false positive forecasts is denoted by FP . They refer to the situations in which an avalanche occurred when there was none predicted and vice versa, respectively. For model assessment, the following quality criteria were applied:

Sensitivity (Probability of Detection):

$$POD = \frac{TP}{TP + FN} \quad (1)$$

Specificity (Probability of Non-Event):

$$PON = \frac{TN}{TN + FP} \quad (2)$$

False Alarm Ratio:

$$FAR = \frac{FP}{FP + TP} \quad (3)$$

True Skill Statistic:

$$TSS = \frac{TP}{TP + FN} - \frac{FP}{FP + TN} \quad (4)$$

Positive Predictive Value:

$$PPV = \frac{TP}{TP + FP} = 1 - FAR \quad (5)$$

Negative Predictive Value:

$$NPV = \frac{TN}{TN + FN} \quad (6)$$

Amongst these, the positive predictive value and the negative predictive value [14] are most informative regarding the operational use of a model used for decision support in our application scenario. Given a negative forecast, the negative predictive value is the probability for this forecast being right. Given a positive forecast, the positive predictive value is the probability for an avalanche event.

The ideal classifier shows a high sensitivity as well as a high specificity. But these quality criteria do not allow for an assessment of the forecast. Obviously, high sensitivity and high specificity result in high positive and negative predictive values.

The probability of detection, the probability of a non-event, the false alarm ratio as well as the true skill statistic are established quality measures for model assessment in avalanche forecasting and are adopted in [12, 17, 18].

4 Data

The Canton of Glarus is located in the eastern part of Switzerland and is characterized by high mountains and steep slopes. In this work, we focused on the alpine valley Kleintal situated in the southeast of the Canton of Glarus. The valley floor of the Kleintal is gently inclined, its elevation ranging from over 600 m.a.s.l. to over 1000 m.a.s.l. The starting zone of a snow avalanche may be situated up to 1700 m above the valley floor and may therefore endanger the main road leading through the valley. The data consist of meteorological variables measured daily in the early morning as well as avalanche information between January 1st, 1972 and April 30th, 2013.

The measures were collected in Elm at 958 m.a.s.l. and at Risiboden, a location situated 2.5 km from Elm at an elevation of 1690 m.a.s.l. They comprised

the maximum and minimum air temperature in the last 24 hours, actual wind speed and actual wind direction, degree of sky cover and precipitation in the last 24 hours in Elm as well as snow depth and new snow depth in the last 24 hours at Risiboden. The air temperature is measured in Celsius degrees and recorded in $1/10$ Celsius degrees in the NXD2000 database. The wind direction is measured in arc degrees and the rounded value of the measured value divided by 10 is recorded with 0 or 36 indicating wind coming from north, and 9, 18, and 27 indicating wind coming from east, south, and west, respectively. The wind speed is measured in meters per seconds and recorded in knots. For standardization purposes, for this work, the wind direction was set to 0 where either the wind speed was 0 or the wind direction was 36. The degree of sky cover was recorded as follows: 0 indicates a clear sky, 4 a coverage of 50% and 8 a cloudy sky. For the precipitation, the water equivalent was given, i.e. the snow was melted and the water amount was recorded in millimeters. The new snow depth as well as the snow depth are measured in centimeters and recorded unaltered.

Meteorological factors are potentially useful for estimating snowpack instability, but interpretation is uncertain and the evidence less direct than for snowpack factors [16]. Avalanche expert knowledge was taken into account by using the derived variables listed in Table 1, which were defined for NXD2000 for the Canton of Glarus and are documented for internal use. In the following, we explain how these variables and their range of values are derived from the data described above.

Table 1. Meteorological variables are measured daily. The definition of derived variables allows to consider an expert knowledge about avalanche activity.

Variable	Unit	Range of values
1 Max. air temperature in the last 24 hours	$[1/10 \text{ } ^\circ\text{C}]$	$[-178, 240]$
2 Max. air temperature in the last 48 to 24 hours	$[1/10 \text{ } ^\circ\text{C}]$	$[-178, 240]$
3 Min. air temperature in the last 24 hours	$[1/10 \text{ } ^\circ\text{C}]$	$[-251, 157]$
4 Min. air temperature in the last 48 to 24 hours	$[1/10 \text{ } ^\circ\text{C}]$	$[-251, 157]$
5 Actual wind direction		$\{0, 10, \dots, 350\}$
6 Wind direction of the previous day		$\{0, 10, \dots, 350\}$
7 Wind speed	[kn]	$[0, 206]$
8 Wind speed of the previous day	[kn]	$[0, 206]$
9 Degree of sky cover		$\{0, 12, \dots, 96\}$
10 Precipitation in the last 24 hours	[mm]	$[0, 989]$
11 Precipitation in the last 48 to 24 hours	[mm]	$[0, 989]$
12 New snow fallen in the last 24 hours	[cm]	$[0, 550]$
13 New snow fallen in the last 72 to 24 hours	[cm]	$[0, 575]$
14 Snow depth		$[0, 432]$

The maximum and minimum air temperature in the last 48 to 24 hours (lines 2 and 4 in Table 1) refers to the maximum and minimum air temperature

recorded for the previous day. The wind direction (lines 5 and 6 in Table 1) is multiplied by 10, the wind speed (lines 7 and 8) is multiplied by 5.1479, and the degree of sky cover (line 9) is multiplied by 12. The precipitation in the last 48 to 24 hours (line 11) refers to the precipitation recorded for the previous day. The amount of new snow fallen in the last 24 hours (line 12) is multiplied by 5. The amount of new snow fallen in the last 72 to 24 hours (line 13) is defined as the sum of the weighted new snow depths of the last 3 days multiplied by 5. The snow depth (line 14) is divided by the mean of all snow depths in the database and multiplied by 100.

Only the avalanches endangering the main road were recorded. In this work, we included 7 avalanche paths with 7 to 13 avalanches each. We did not discriminate between avalanche paths, and days with at least one avalanche being released in one of these paths were considered as one event. The complete data set contained 53 positive examples, i.e. avalanche days, and 6889 negative examples, i.e. non-avalanche days. The ratio of positive to negative examples therefore was approximately 1:130. The data set was divided into a training and a test set as follows: The test set consisted of all entries from November 1st, 2002 to April 30th, 2013. By this means, the ratio of positive to negative examples matched approximately the ratio observed in the real world. The training set consisted of all avalanche days from January 1st, 1972 to April 30th, 2002 and about 10 times as many non-avalanche days drawn randomly every year. The test data set consisted of 12 positive and 1572 negative examples, the training data set consisted of 41 positive and 560 negative examples.

5 Results

Two BRFs were trained using the size of the positive class as bootstrap sample size for positive and negative examples. For the positive class, the cutoffs were set to 0.5 and 0.6, respectively. The number of variables to be tested for the best split was set to 2. The cutoff as well as the number of variables to be tested for the evaluation of the splits were determined using 10-fold cross-validation. For the cutoff, the following values were tested: 0.3, 0.4, 0.5, 0.6, and 0.7. For the number of variables, the following values were tested: 2, 3, and 4. The contingency tables obtained for the test data are shown in Fig. 2.

		Observed				Observed			
			0	1				0	1
Predicted	0	1496	6	Predicted	0	1517	7		
	1	76	6		1	55	5		
Model BRF_0.5				Model BRF_0.6					

Fig. 2. For model BRF_0.5, the cutoff for the positive class was set to 0.5. For model BRF_0.6, the cutoff for the positive class was set to 0.6. It is observed that a decrease in the cutoff leads to an increase in the number of true positives and false positives.

For model BRF_0.5 with a cutoff of 0.5, the number of true positives is higher with respect to model BRF_0.6 with a cutoff of 0.6. On the other hand, model BRF_0.6 achieved a lower number of false positives. Using WRF, model WRF_5 was trained using a class weight of 5 for the positive class and a class weight of 1 for the negative class. The class weights were determined using 10-fold cross-validation. For the positive class, the following weights were tested: 1, 2, 3, 5, 10, 15, 20, 25, 30, 50, 100, 110, 120, 130, and 150. The class weight for the negative class was set to 1. The contingency table obtained for the test data is shown in Fig. 3. In Table 2, the quality measures for the generated models for the test data are listed.

		Observed	
		0	1
Predicted	0	1516	7
	1	56	5

Model WRF_5

Fig. 3. For model WRF_5, the class weights for the positive and negative class were set to 5 and 1, respectively. The results are very similar to the ones obtained for model BRF_0.6.

Table 2. In the quality measures obtained for the test data, the analogy between models BRF_0.6 and WRF_5 becomes evident. No WRF similar to model BRF_0.5 could be trained with an acceptable number of false positive forecasts.

Model	TN	FN	FP	TP	POD	PON	FAR	TSS	PPV	NPV
BRF_0.5	1496	6	76	6	0.500	0.952	0.927	0.452	0.073	0.996
BRF_0.6	1517	7	55	5	0.417	0.965	0.917	0.382	0.083	0.995
WRF_5	1516	7	56	5	0.417	0.964	0.918	0.381	0.082	0.995

The performance of models BRF_0.6 and WRF_5 was almost the same. The sensitivity was quite low with only 41.7%, hence the number of false positives was low too compared to model BRF_0.5. No WRF with a sensitivity of 50% and an acceptable number of false positive forecasts could be trained. While in the case of a positive forecast an avalanche occurred only with a probability of 7.3% to 8.3%, depending on the model used, negative forecasts were very reliable with a negative predictive value of 99.5% and 99.6%, respectively. Compared to the quality measures obtained for the training data shown in Table 3, the sensitivity, the true skill score and the positive predictive value were considerably lower.

Table 3. The quality measures obtained for the training data showed differences between BRF and WRF. The sensitivity for the BRFs was 100%, while for model WRF_5 it was 65.9%.

Model	TN	FN	FP	TP	POD	PON	FAR	TSS	PPV	NPV
BRF_0.5	525	0	35	41	1.000	0.938	0.461	0.938	0.539	1.000
BRF_0.6	538	0	22	41	1.000	0.961	0.349	0.961	0.651	1.000
WRF_5	531	14	29	27	0.659	0.948	0.518	0.607	0.482	0.974

The BRF models showed a sensitivity of 100% for the training data while for model WRF_5, a sensitivity of 65.9% was obtained. For all models, the sensitivity, the true skill statistic and the positive predictive value were noticeably higher for the training data than for the test data. Accordingly, the false alarm ratio was lower for the training data than for the test data. These differences were more pronounced with the BRF models respect to the WRF model.

The similarity of the two BRF models became visible in the misclassified examples: For the test data, the false negatives for model BRF_0.5 were a subset of the false negatives for model BRF_0.6. On the other hand, the false positives for model BRF_0.6 were a subset of the false positives of model BRF_0.5. When comparing the misclassified examples for the BRF and WRF models, it was noticed that both models showed the same false negative predictions. 46 false positives showed up in BRF_0.6 as well as in WRF_5. The comparison of all three models can be summarized as follows: 52 examples were misclassified by all three models, consisting of 6 false negatives and 46 false positives. This makes up 63.4% of the misclassifications of model BRF_0.5, 83.9% of the misclassifications of model BRF_0.6 and 82.5% of the misclassifications of model WRF_5.

Compared to the two decision trees described in [12], our models achieved a lower sensitivity but a higher specificity. With different test sets, the models described in [12] had a sensitivity of 61% and 100% with a corresponding specificity of 83% and 21%, respectively. The true skill score was 21% for the second model and therefore lower than in our models. The first model achieved a true skill score of 44% which was lower than in our model BRF_0.5 only. The false alarm ratio in our models is noticeably higher for the test data. With the training data, our models achieved a higher true skill score, specificity and false alarm ratio than did the models in [12]. It has to be remarked that in [12] the ratio of positive to negative examples in the test data was about 1:6. Furthermore, the presented models were trained considering two or three variables while in our models we applied no variable selection.

In [18], large and infrequent snow avalanches are predicted. Considering new snow depth only, a sensitivity of about 65% was achieved which seems favourable compared to our models. The presented models showed a false alarm ration of about 90% which is similar to the false alarm ration of our models on the test data.

For the purpose of comparison, the following classifiers were employed using the default settings in WEKA [11]: AdaBoost.M1 [8] using DecisionStump [13]

as base classifier; bagging [1] using DecisionStump [13] and REPTree [7] as base classifier; logistic regression [15]. The resulting models show a significantly lower sensitivity compared to our models and therefore are not applicable for decision support in our case.

6 Discussion

Two types of models proved to be feasible: On the one hand, two models with a sensitivity of slightly more than 40% were trained. On the other hand, one model with a sensitivity of 50% was trained. The latter showed a considerably higher number of false positive forecasts and a slightly lower positive predictive value. No WRF with a sensitivity of 50% and an acceptable number of false positives could be trained. Therefore, with this data, BRF could be chosen for modeling a system for decision support in avalanche warning.

The trained models are feasible as a decision support in avalanche warning: The testing period comprises 11 winter seasons consisting of approximately 181 days each. For this period, 55 to 76 false positive forecasts are acceptable. The misclassification rate is comparable to that of an human expert. Accepting a higher number of false negative forecasts for a higher sensitivity may make sense: Not all avalanches contained in the database reached the road. Danger does not occur always with a false negative forecast and therefore 6 to 7 false negatives are acceptable. It would be interesting to differentiate between avalanches that reached the road and avalanches that did not. In this work, due to the lack of data, this differentiation was not made.

The models are developed for the Kleintal in the Canton of Glarus in Switzerland based on the data contained in NXD2000. In contrast to NXD2000, no comparison with previous similar situations can be made, but the models allow for probabilistic forecasts. Therefore, the cutoff for the positive class for BRF could also be determined using ROC curves and the corresponding weight adopted in WRF for the positive class derived from this cutoff according to the procedure described in [6].

The positive and negative predictive values present valuable information for assessing a given forecast. While sensitivity and specificity are important quality measures and their values have to be as high as possible, they do not allow for an assessment of a given prediction. However, this is an important information for the avalanche service using the system as a decision support in avalanche warning.

In BRF_0.5, BRF_0.6 and WRF_5, mostly the same examples were misclassified. Considering the fact that particularly for the negative examples the misclassified examples comprised less than 5% of all examples, it can be supposed that these misclassifications are due to the data. The training examples were chosen randomly on a yearly basis and therefore few consecutive days are present in the training data set. The training examples could as well be chosen using statistical methods analogous to the approaches employed in [12, 17, 18]. The test data set contains time series of meteorological variables for up to 178 consecu-

tive days. Differentiation of two consecutive days belonging to different classes poses a major challenge and cannot be made by analyzing the meteorological values only. The definition of additional meaningful variables could improve the differentiation between positive and negative examples.

The model performance with the training data is significantly superior to the model performance with the test data. Generalization seems to be an issue. One possible reason may be that the test data contains a high percentage of consecutive days, thus the recommendations given in the last paragraph apply.

7 Conclusions and Further Work

Based on meteorological data measured on a daily basis as well as avalanche data, a system for the decision support in avalanche warning has been modeled. In this data, avalanche days are rare events. All trained models have a maximum sensitivity of 50% and a high false alarm ratio. Nevertheless, the trained models are feasible as decision support in avalanche warning. The number of false negative and false positive forecasts are acceptable with respect to the period considered, and approximately match the performance of a human expert.

Compared to the models described in [12, 17, 18], the following aspects of our work should be noted: First, the quality measures were chosen from the point of view of the designated user, the members of the avalanche service which is in charge of assessing the local avalanche danger. The positive and negative predictive values, which are not presented in the cited approaches, provide an assessment of the forecast reliability. From an operational point of view, these are the most important quality measures. Second, the models allow for probabilistic forecasts and therefore for the characterization of the probability of an event. Third, BRF and WRF proved an adequate starting point for obtaining a feasible system for decision support in snow avalanche warning with rare events.

There are several directions in which the work presented in this paper should be extended. In order to achieve a higher performance, additional meaningful variables should be defined; these may be quantitative as well as qualitative variables. A more sophisticated variable selection could also prove beneficial. Depending on the weather situation, the importance of meteorological variables varies. This can be taken into account by defining variables describing the weather situation. Since the snowpack develops with time, defining variables characterizing weather trends could prove advantageous. The influence of the training data on the generated model should be investigated. Training examples may be chosen according to statistical criteria, or all data not assigned to the test data set may be used for training. The possibility to additionally predict which avalanche path is in danger of being released would be advantageous for the avalanche service. However, this requires an appropriate amount of data.

References

1. Breiman, L.: Bagging predictors. *Machine Learning* 24(2), 123–140 (1996)

2. Breiman, L.: Random Forests. *Machine Learning* 45(1), 5–32 (2001)
3. Breiman, L., Friedman, J., Stone, C.J., Olshen, R.A.: *Classification and regression trees*. CRC Press (1984)
4. Buser, O., Bütler, M., Good, W.: Avalanche forecast by the nearest neighbour method. *International Association of Hydrological Sciences* 162, 557–570 (1987)
5. Chen, C., Liaw, A., Breiman, L.: Using Random Forest to Learn Imbalanced Data. Tech. rep., University of California, Berkeley (2004), <http://www.stat-www.berkeley.edu/tech-reports/666.pdf>
6. Elkan, C.: The foundations of cost-sensitive learning. In: *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*. pp. 973–978 (2001)
7. Elomaa, T., Kääriäinen, M.: An Analysis of Reduced Error Pruning. *Journal of Artificial Intelligence Research* 15, 163–187 (2001)
8. Freund, Y., Schapire, R.E.: Experiments with a New Boosting Algorithm. In: *International Conference on Machine Learning*. vol. 96, pp. 148–156 (1996)
9. Gassner, M., Brabec, B.: Nearest neighbour models for local and regional avalanche forecasting. *Natural Hazards and Earth System Sciences* 2, 247–253 (2002)
10. Gassner, M., Etter, H.J., Birkeland, K., Leonard, T.: NXD2000: An improved avalanche forecasting program based on the nearest neighbor method. In: *Proceedings of the 2000 International Snow Science Workshop*. pp. 52–59 (2000)
11. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA Data Mining Software: An Update. *SIGKDD Explorations* 11(1) (2009)
12. Hendrikx, J., Murphy, M., Onslow, T.: Classification trees as a tool for operational avalanche forecasting on the Seward Highway, Alaska. *Cold Regions Science and Technology* 97, 113–120 (2014)
13. Iba, W., Langley, P.: Induction of One-Level Decision Trees. In: *Proceedings of the Ninth International Conference on Machine Learning*. pp. 233–240 (1992)
14. Lalkhen, A.G., McCluskey, A.: Clinical tests: sensitivity and specificity. *Continuing Education in Anaesthesia, Critical Care & Pain* 8(6), 221–223 (2008)
15. Le Cessie, S., Van Houwelingen, J.: Ridge Estimators in Logistic Regression. *Applied Statistics* 41(1), 191–201 (1992)
16. McClung, D., Schaerer, P.A.: *The Avalanche Handbook*. The Mountaineers Books (2006)
17. Mitterer, C., Schweizer, J.: Analyzing the atmosphere-snow energy balance for wet-snow avalanche prediction. In: *Proceedings of the International Snow Science Workshop*. pp. 77–83 (September 2012)
18. Schweizer, J., Mitterer, C., Stoffel, L.: On forecasting large and infrequent snow avalanches. *Cold Regions Science and Technology* 59(2), 234–241 (2009)
19. Weiss, G.M.: Mining with Rarity: A Unifying Framework. *ACM SIGKDD Explorations Newsletter* 6(1), 7–19 (2004)
20. WSL-Institut für Schnee- und Lawinenforschung SLF, Bundesamt für Umwelt BAFU, Schweizerische Interessengemeinschaft Lawinenwarnsysteme (SILS): *Praxishilfe. Arbeit im Lawindienst: Organisation, Beurteilung lokale Gefährdung und Dokumentation*. http://www.slf.ch/dienstleistungen/merkblaetter/praxishilfe_lawdienst_deutsch.pdf, information sheet, (in German)