Cold Regions Science and Technology 111 (2015) 60–66

Contents lists available at ScienceDirect

Cold Regions Science and Technology

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

Robust snow avalanche detection using supervised machine learning with infrasonic sensor arrays

Thomas Thüring *, Marcel Schoch, Alec van Herwijnen, Jürg Schweizer

WSL Institute for Snow and Avalanche Research SLF, Davos, Switzerland

ARTICLE INFO

Article history:
Received 2 July 2014
Received in revised form 3 November 2014
Accepted 26 December 2014
Available online 4 January 2015

Keywords:
Snow avalanches
Infrasound
Array processing
Automated monitoring
Machine learning

ABSTRACT

Automated detection of snow avalanches is crucial to assess the effectiveness of avalanche control by explosions, and to monitor avalanche activity in a given area in view of avalanche forecasting. Several automated or semi-automated detection technologies have been developed in the past among which infrasound-based detection is the most promising for regional-scale avalanche monitoring. However, due to significant ambient noise content in infrasonic signals, e.g. from atmospheric processes or airplanes, fully automated and reliable avalanche detection has been very challenging. Signal processing is highly critical and strongly affects detection accuracy. Here, a robust detection method by using supervised machine learning is introduced. Machine learning algorithms can take into account multiple signal features and statistically optimize the classification task. We analyzed infrasound data with concurrent visual avalanche observations from the test site Lavin (Eastern Swiss Alps) for the winter of 2011–2012. A support vector machine was trained by using training data from the first half of the winter season and the accuracy was tested on data from the second half of the season. A significant reduction of false detections, from 65% to 10%, was achieved compared to a threshold-based classifier provided by the sensor manufacturer. The proposed method enables reliable assessment of the avalanche activity in the surroundings of the system and paves the way towards robust and fully automated avalanche detection using infrasonic systems.

1. Introduction

Snow avalanches threaten people and infrastructure in seasonally snow-covered mountain regions. Thorough avalanche hazard assessment is crucial for minimizing the risk by avalanche preventive measures such as avalanche control by artificial triggering of avalanches. Hazard assessment relies on weather data and forecasts, snow cover model output and snow instability data (McClung and Schaerer, 2006; Schweizer and Jamieson, 2010). The latter includes field observations of snow instability, in particular results of snowpack stability tests and avalanche occurrence data. Complete assessment of avalanche activity in an area requires continuous monitoring, which cannot be achieved with field observations. In particular during times of poor visibility or at night, when field observations are impossible, automated detection systems are highly desirable. A fully automated system continuously observes an area and generates events which are transmitted to the avalanche safety service in charge. These systems are also needed at avalanche control sites to measure the effectiveness of the artificial triggering by explosions (Schweizer and van Herwijnen, 2013).

A variety of remote sensing techniques and instruments for the automated detection of snow avalanches have been reported in the past. Technologically, they can be classified into techniques based on radio frequency signals (radars) (Gauer et al., 2007; Kogelnig et al., 2012; Salm and Gubler, 1985; Vriend et al., 2013), seismic signals (geophones) (Schaerer and Salway, 1980; Suriñach et al., 2000; van Herwijnen and Schweizer, 2011a, 2011b), optical signals (imagery) (Larsen et al., 2010; Lato et al., 2012) and acoustic signals (microphones, micro barometers) (Adam et al., 1998; Bedard, 1989; Kogelnig et al., 2011; Scott et al., 2007; Ulivieri et al., 2011). Optical imagery enables the assessment of avalanche activity and localization with high spatial resolution; however, its applicability strongly depends on visibility. Avalanche detection using pulsed Doppler radar is very reliable and enables the measurement of avalanche dynamics properties, for instance avalanche velocity. On the other hand, the monitoring area is usually small, typically a well-defined avalanche path. Seismic detection performs well during all weather conditions. However, seismic signals from natural or artificial sources (e.g., earthquakes, airplanes) cause significant background noise (van Herwijnen and Schweizer, 2011b). Several automated detection approaches to separate noise from avalanche events in seismic data have been reported in the past (Bessason et al., 2007; Lacroix et al., 2012; Leprettre et al., 1996).

Similar to seismic waves, flowing and turbulent snow masses by avalanches generate pressure waves (or sound waves) in the air.

* Corresponding author at: Unterhunzikon 1, 6232 Geuensee, Switzerland. Tel.: +41 797770641.
E-mail address: t.thuering@gmx.net (T. Thüring).

http://dx.doi.org/10.1016/j.coldregions.2014.12.014
0165-232X/

These sound waves lie in the frequency range between 0.001 and 20 Hz, known as the infrasonic range (Bedard, 1989). The atmospheric attenuation of infrasound is very low and the waves can travel over large distances. Therefore, infrasonic detection systems can potentially record signals many kilometers away. Infrasound has already been used for monitoring atmospheric processes (Le Pichon et al., 2009), volcanic activity (Ripepe et al., 2007), nuclear explosions (Christie et al., 2001) or snow avalanches (Scott et al., 2006). In mountain regions, infrasonic sensors enable the detection of avalanches releasing several kilometers away from the sensor system (Ulivieri et al., 2011), which is a major benefit compared to other detection techniques mentioned above. Infrasonic systems may consist of multiple sensors (arrays), which enhance the signal to noise ratio (SNR) (Scott et al., 2007) and enable precise source localization (Ulivieri et al., 2011). Several commercial products for automated avalanche detection have been released recently, indicating an increasing demand of infrasound as an alternative or complementary technique to existing technologies. Examples are the ARFANG system by IAV Switzerland or the system by iTEM geophysics in Italy.

Two major issues in infrasonic-based avalanche detection are the presence of ambient noise (e.g., from wind) and of signal sources other than avalanches (e.g. atmospheric processes, airplanes, helicopters, etc.). If such disturbing signals cannot efficiently be separated from avalanche signals, false detections (i.e. false alarms) may frequently occur. Ambient noise can be reduced, either by adding noise filters or by using multiple, spatially distributed sensors (arrays) (Scott et al., 2007). Noise signals of a sensor array are mutually uncorrelated for a sufficiently high sensor spacing, which can be used to increase the overall SNR. Other than uncorrelated noise, correlated signals from real infrasonic sources can become a serious problem, as they may be difficult to separate from avalanche signals. To overcome this, robust signal analysis and classification methods are required. Several classification approaches have been proposed (Chritin et al., 1996; Schimmel and Hübl, 2013; Ulivieri et al., 2011). These methods all use a common classification scheme. First, signal features, such as the infrasonic power, the direction of incidence of an event or the duration of an event, are extracted. Second, features are analyzed and thresholds are defined which separate avalanche events from non-events. Such threshold-based classifiers are easy to control and may perform well under certain conditions. Unfortunately, classification accuracy has not been reported in these studies.

Here, a threshold-based classifier, developed and optimized by the manufacturer of a commercial infrasonic avalanche detection system, is compared to an alternative, machine learning-based approach. A support vector machine (SVM), a well-established machine learning algorithm, was trained and evaluated using infrasonic recordings from a four-sensor-array system installed at an avalanche control site near Lavin in the Eastern Swiss Alps. SVM-based avalanche detection was already demonstrated for the automated detection of avalanches in seismic data, showing encouraging results (Rubin et al., 2012). Using an infrasonic system, an SVM was also applied to the detection of volcanic activities at Mount Etna (Cannata et al., 2011).

In Section 2, after a short overview of the infrasonic sensor hardware, the computational methods are discussed in detail. While SVMs are nowadays a standard tool in data mining problems and can easily be applied to avalanche detection tasks, the extraction of discriminant features which separate positive from negative events remains a crucial and time consuming issue. In Section 3, an SVM is trained based on various signal features and by using avalanche field observation data in the avalanche controlled area from the early winter season 2011–2012 (training phase). Detection performance in the avalanche controlled area is evaluated by using a test data set from the remaining season in 2012 (test phase). Finally, avalanche activity is assessed for the area in the vicinity of the avalanche path during the test phase.

2. Instrumentation and methods

2.1. Infrasonic sensors

A commercial infrasonic sensor array with four sensors from IAV Engineering (Tannay, Switzerland) was installed in 2009 near Lavin (Eastern Swiss Alps) to monitor the Gonda avalanche path where avalanches are triggered artificially by explosives to protect the road passing below. The system is located across the valley from the avalanche path at the bottom of the counter slope; for more details on the site see Meier and Lussi (2010). The four sensors are aligned in a star-shaped geometry with a radius of 30 m to equalize the angular reception sensitivity (Ulivieri et al., 2011; Van Lancker, 2001). The sensors measure pressure variations differentially with respect to the atmosphere (microbarometer) and synchronously acquire infrasonic signals at a sampling rate of 80 Hz. With a cutoff frequency of approximately 0.1 Hz, the available frequency spectrum ranges from 0.1 to 40 Hz. The sensitivity of the sensors is 3.2 V/Pa, the dynamic range is 80 dB and the noise floor is −85 dBFS. Compared to a single sensor, a sensor array has two main advantages. First, uncorrelated noise is suppressed due to the spatial distribution of the sensors. Second, it enables localization of signal sources by measuring time delays of signals between sensor pairs (see Section 2.3). The resolution of source localization depends on the direction of the moving source, the sensor geometry and the sampling frequency. For the system used here, the maximum achievable angular resolution is $\alpha_{\min} \approx 4°$.

2.2. Event classification by supervised learning

Event classification from acoustic signals is a wide research area with a large variety of available algorithms and techniques. The common goal is to classify recorded signals into events based on a set of extracted variables (features) from the raw output signal. A feature represents a signal property which can ideally separate the signal into different classes of interest. A simple and common scheme of a classifier is to define decision thresholds for the extracted features. This usually allows for simple control and optimization of the classification outcome. However, if multiple features must be taken into account, manual definition and optimization of thresholds become difficult to handle and are often prone to mistakes (Kotsiantis et al., 2007). In particular when trying to establish relationships between features, threshold-based classification is not practical.

In contrast to threshold-based classification, machine learning algorithms automate and statistically optimize a classification task (Vapnik, 2000). Supervised machine learning methods “learn” decision margins of features from training data. Therefore, a set of training data with known classification output must be available a-priori. Among the large variety of supervised machine learning algorithms, support vector machines (SVM) are known to perform well when using multiple and continuous features (Kotsiantis, 2007). The algorithm analyzes the multidimensional feature space and calculates optimum decision margins (hyper planes) based on the training data. This allows for taking into account complex mutual dependencies between features, which is a major benefit compared to threshold-based methods. Only a few input parameters are required. A detailed discussion about SVMs and classification theory can be found in the literature (Burges, 1998; Cristianini and Shawe-Taylor, 2000).

We applied an SVM to the classification problem of detecting avalanche events in infrasonic signals. Fig. 1 shows a schematic overview of the different steps in the proposed signal processing workflow. The method is implemented “off-line”, meaning that previously recorded infrasonic data are analyzed. The raw data were preprocessed before extraction of signal features. After feature extraction, part of the data was assigned as training data, which were used to train the SVM classifier (learning phase). The classifier was optimized through a 10-fold cross-validation procedure. Finally, the detection of avalanche

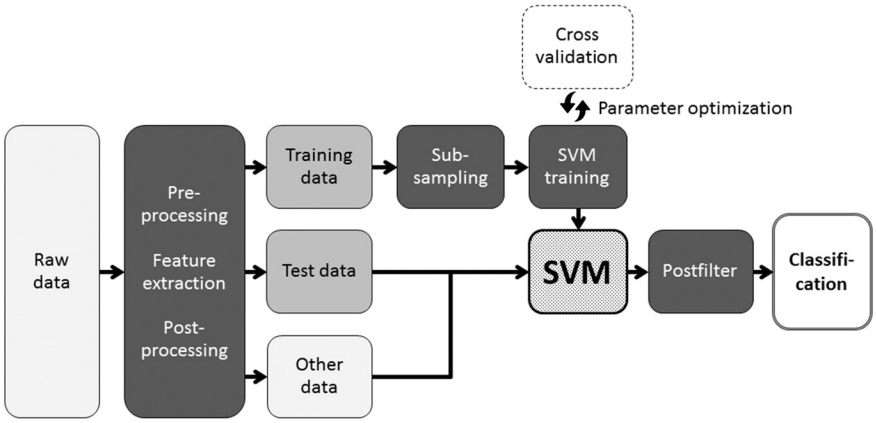

Fig. 1. Signal processing workflow for avalanche classification based on an SVM. After preprocessing the raw data and the feature extraction, the data are split into three groups: training data, test data and other data. Training and test data are labeled, which means that the classification is known a-priori from avalanche observation data. In order to obtain the same number of positive and negative event vectors for the training of the SVM, training data must usually undergo some subsampling. The SVM performance is optimized by minimizing the error rate obtained from a 10-fold cross validation. With the trained SVM, test data and other data can be classified. The last stage, a post filter, is used to render the SVM classification output into avalanche events according to fixed criteria, for instance by avalanche duration or by azimuth angle.

events was obtained by applying a post processing filter to the classification output of the SVM.

The following paragraphs discuss the details of the implemented classification scheme and the associated signal processing workflow of Fig. 1 in more detail.

2.3. Preprocessing and feature extraction

Signal preprocessing involves preparation of the raw sensor signals $s_i(t)$ with $i \in [1, N]$ (sensor index with N the number of sensors) for the subsequent feature extraction. In a first step, by applying a moving window of length $t_w = 4\ s$ and overlap $t_o = 3\ s$ (75%) to $s_i(t)$, n_f signal frames $f_{i,j}(t)$ ($j \in [1, n_f]$) are generated. In continuous recording mode, this would correspond to the extraction of a 4-second signal frame which is updated every second ($\Delta t_f = t_w - t_o = 1$ s). Features are extracted based on these signal frames, which implies that the feature vector (vector containing all features for one frame) is updated every second and thus classification can be performed at this rate.

The next step is the extraction of discriminant signal features on each frame. Feature extraction is a key step as the accuracy of any classification task is strongly dependent on the careful selection and generation of discriminant features. They must efficiently separate positive events from negative events. Here, all features are time dependent and are updated at (Δt_f) and the extraction of a feature in a signal frame mainly involves three steps: pre-filtering of the signal frames, calculation of the feature and post-filtering of the time-dependent feature signal. While pre-filtering aims at optimum signal preparation (high SNR) for a specific feature, the post filter is mainly used for general de-noising. Table 1 shows an overview about the extracted features from the signals and the applied pre and post-filters. In the following, details of the feature extraction step are discussed.

One class of features is the Fourier spectrum-based features. These features require the computation of the power spectrum from the signal frames, given by

$$S_{i,j}(\nu) = \left|\mathrm{FT}\left(f_{i,j}(t)\right)\right|^2, \tag{1}$$

where ν is the frequency and FT(□) refers to the Fourier transform. Using multiple sensors, the Fourier frames $S_{i,j}(\nu)$ can be multiplied, yielding

$$S_j(\nu) \equiv S(\nu) = \prod_{i=1}^{N} S_{i,j}(\nu). \tag{2}$$

The multiplied power spectrum, similar as the cross-correlation of two signals, suppresses uncorrelated signals and enhances correlated signals between the sensors and thus exhibits an increased SNR compared to the power spectra $S_{i,j}(\nu)$ of the individual sensors.

The features “low frequency band” and “high frequency band” are measurements of the signal energy content in the low and high-frequency band of $S(\nu)$, respectively. For the low-frequency band, it was found that the energy content between 6 and 8 Hz, measured relative to the total signal energy, provides a discriminant feature.

Table 1
Features extracted for the training of the SVM. Each feature extraction involves preprocessing, extraction and post processing steps. Features are calculated at the frame frequency, which is $\Delta t_f = 1\ s^{-1}$. The last three features in the table are standard deviation (STD) or mean values of previous features, measured during a time period of 13 s. Post processing usually consists of a median filter for de-noising. The right column further shows the minimum error rate obtained from 10-fold cross-validation for the classification using only one feature. This indicates the power of each feature to separate the data into positive and negative events. Using all features for an SVM classification, the cross-validation error rate was 8.6%.

Feature	Preprocessing	Feature calculation	Post processing	CV error rate
Low frequency band	See Eq. (2)	Frequency band: 6 to 8 Hz	Median filter: 5 s	35.8%
High frequency band	See Eq. (2)	Frequency band: 26 to 35 Hz	Median filter: 5 s	23.4%
Frequency band ratio	See Eq. (2)	7 to 16 Hz relative to 16 to 37 Hz	Median filter: 5 s	24.4%
Mean frequency ($\bar{\nu}$)	None	See Eq. (3)	Median filter: 5 s	44.4%
Frequency variance (M_2)	None	See Eq. (4)	Median filter: 5 s	41.1%
Frequency kurtosis (M_4)	None	See Eq. (4)	Median filter: 5 s	35.9%
Mean time residual	Time delay estimation	Sum of delay times in each sensor triplet; mean of all time residuals.	None	30.3%
Azimuth angle velocity	Source localization	Derivative of azimuth angle	Median filter: 5 s	17.4%
Elevation angle velocity	Source localization	Derivative of elevation angle	Median filter: 5 s	22.1%
STD azimuth angle	Source localization	STD during 13 s	None	12.3%
Mean azimuth angle	Source localization	Mean during 13 s	None	11.3%
STD azimuth velocity	Source localization	STD during 13 s	None	11.6%

No pre-filter is applied and the post filter is a median filter of 5 second duration. The high-frequency-band feature extracts the signal energy between 26 and 35 Hz. The feature "frequency band ratio" measures the ratio of the energy content in the frequency bands from 7 to 16 Hz and 16 to 37 Hz.

The feature "mean frequency" is given by

$$\overline{\nu} = \int_0^{\nu_N} \nu\, S(\nu)\, d\nu. \tag{3}$$

The features "frequency variance" and "kurtosis" are based on the calculation of the 2nd (n = 2) and 4th (n = 4) central moments, respectively, of the frequency spectrum, given by

$$M_n = \int_0^{\nu_N} (\nu - \overline{\nu})^n\, S(\nu)\, d\nu. \tag{4}$$

Another class of features is the source localization-based features. Source localization aims at retrieving the position of a sound source by using the delays of the signal arrival times in a sensor array (Johnson and Dudgeon, 1992). The first step is the calculation of the signal delay times for each sensor pair, a process known as time delay estimation (TDE). Robust TDE has extensively been studied for a variety of applications. Here, a standard approach is used, which is based on the pairwise cross-correlation of the sensor signals (Knapp and Carter, 1976). Prior to cross-correlation, the signal can be filtered to increase the signal-to-noise ratio (SNR) and improve TDE for specific signal sources such as avalanches. However, the increased SNR usually comes at the expense of a reduced time delay resolution. Here, a Gaussian low-pass filter with a mean of 7 Hz and a full-width-half-maximum (FWHM) of 2 Hz was applied to the signal frames.

The time residual is a feature which is directly extracted from the TDE. It is defined as the sum of delay times between each pair of sensors in a closed loop. For a parallel wave passing through the array with constant velocity, the time residual must be zero. To form a loop, at least 3 sensors are required and the time residual in such a sensor triplet can be expressed by

$$t_{res} = \tau_{1,2} + \tau_{2,3} + \tau_{3,1}, \tag{5}$$

where $\tau_{i,j}$ is the time delay from sensor i to sensor j (Ulivieri et al., 2011). With an array of N sensors, the number of triplets can be obtained from the binomial coefficient, given by

$$n_t = \binom{N}{3} = \frac{N!}{3!(N-3)!}. \tag{6}$$

For a four-sensor array, $n_t = 4$. The feature "mean time residual" in Table 1 is calculated by taking the mean value of the four time residuals.

The second step of source localization is the computation of the source back-azimuth angle by using the calculated time delays and the array geometry. The back-azimuth angle, hereafter referred to as azimuth angle, is the direction of the sound wave incident to the sensor-array with respect to a vector pointing north (Ulivieri et al., 2011). Again, several approaches for the calculation of the azimuth angle exist which mainly depend on the number of sensors used. A four-sensor array theoretically enables source localization in full 3D space (Van Lancker, 2001). Since the array is operating in the far-field (distance to source much larger than the wavelength), a parallel wave approach was used for the avalanche localization, enabling the calculation of azimuth and elevation angles, of the incoming parallel sound wave (Ulivieri et al., 2011; Van Lancker, 2001). The vector $\vec{n}$ pointing in the direction of the incoming parallel wave was obtained by solving the equation

$$\vec{\tau} c = D\vec{n}, \tag{7}$$

where $\vec{\tau}$ contains the time delay estimates of each sensor pair, D is the sensor-geometry matrix and $c \approx 330$ m/s is the speed of sound in air. Eq. (7) is a linear system of equations with 3 unknowns (components of $\vec{n}$) and $N(N-1)/2$ equations if N is the number of sensors. Having three sensors, D is a square matrix and there is a unique solution to Eq. (7) given by

$$\vec{n} = cD^{-1}\vec{\tau}. \tag{8}$$

Having four or more sensors, Eq. (7) becomes over-determined and the least squares solution can be used, which is given by

$$\vec{n} = c\left(D^T D\right)^{-1} D^T \vec{\tau} = cD^{\dagger}\vec{\tau}, \tag{9}$$

where $D^{\dagger} = (D^T D)^{-1} D^T$ is known as the pseudo inverse of D.

Azimuth (α) and elevation (γ) angle of the wave vector can then be calculated by

$$\alpha = \tan^{-1}\left(\frac{n_y}{n_x}\right), \tag{10}$$

$$\gamma = \sin^{-1}\left(\frac{n_z}{|\vec{n}|}\right). \tag{11}$$

Azimuth and elevation angle provide localization information about the origin of a sound source. In particular when assessing continuously in time, they allow for deriving kinematic properties of a moving source. The features "azimuth angle velocity" and "elevation angle velocity", which are simply time derivatives of Eqs. (10) and (11), provide highly discriminant features for the avalanche detection (see Table 1). Manual inspection of the angle velocities allows for clearly separating avalanche events from non-events. During an avalanche, the azimuth angle and elevation angle velocities usually drop to small (but non-zero) constant values, whereas during non-events, they are completely random indicating that there is no specific source that can be localized.

Features were extracted from each signal frame and are therefore updated at a rate of $\Delta t_f = 1$ s. In order to account for the temporal behavior of features at a longer time scale, some features were further analyzed over a certain period in time. For instance, in addition to the azimuth velocity ($d\alpha/dt$), its standard deviation over 13 s was also used as a feature. An instantly small azimuth velocity is a necessary but not a sufficient condition for an avalanche event. A small standard deviation over several seconds implies a small azimuth velocity over a longer time. The combination of both features enables detecting events which take a certain amount of time, i.e., at least 13 s. The value of 13 s is a result of optimizing the classification rate, which is further discussed in the following paragraph.

Feature extraction introduced additional parameters such as frequency bandwidths or filter lengths, which have an influence on the classification error and eventually on the avalanche detection accuracy and must thus be selected carefully. Here, each parameter of a feature extraction process was determined by minimizing the classification error rate using a classifier which uses only the corresponding feature. The right column of Table 1 shows the minimum achieved error rate for each feature upon sweeping parameter values. For instance, the "low frequency band" feature exhibited a minimum classification error of 35.8% for the frequency band between 6 and 8 Hz. Optimization was performed by sweeping the lower and upper boundaries of the frequency band and by calculating the corresponding classification error. However, the retrieved optima in Table 1 do not provide a general reference for optimum parameters for any infrasonic avalanche detection system, since these parameters probably dependent on multiple factors such as the system response, ambient noise conditions, number of sensors,

distance to avalanche path and so on. On the other hand, the optimization method is generic and can be used on any system.

2.4. Training, cross-validation and testing

Supervised machine learning algorithms require training data to derive the classification margins in the multi-dimensional feature space. Training data, which contains positive and negative events (avalanches and no avalanches), must be labeled, which means that the classification of each feature vector in the training data set must be known a-priori. This information can be obtained either from other (reliable) detection systems or from avalanche observation records. In this study, we used reference data of avalanche events, either spontaneously releasing or artificially triggered and recorded by the local avalanche safety service of the Highways Department by visual inspection. For the training of the SVM, the training data set should contain the same amount of positive events as negative events. However, the number of negative events in the data is usually significantly higher than the number of positive events, that is, the occurrence of avalanche events is rare. Therefore, negative events must be subsampled. In Section 3.1, the subsampling strategy is discussed in more detail.

A linear SVM was trained by solving the dual problem including a box constraint for soft margins and by using the sequential minimum optimization (SMO) algorithm (Cristianini and Shawe-Taylor, 2000). 10-fold cross-validation was used to evaluate and optimize the classification error rate for the training data. The error rate provides a metric on how well the SVM can separate positive from negative feature vectors. It is important to point out that the cross-validation error rate is related to the detection accuracy of feature vectors and is not equal to the detection accuracy of avalanche events in the training-data set. Since the feature vector is updated every second, avalanche events are expected to consist of a series (clusters) of positively classified feature vectors, which must eventually be identified. For this purpose, a post filter can be applied to the classification output. This filter must maintain clusters of positive classifications and suppress single isolated positive classifications. In principle, this can be viewed as a filter for the avalanche duration time. A median filter of length 20 s was used for this task. Finally, a test data set, which is also labeled, enables the assessment of avalanche detection accuracy of the proposed method.

3. Results

3.1. SVM training and cross-validation

Infrasonic data from the winter season 2011/2012 was evaluated for this study. At the avalanche control site, a total of 59 avalanches were observed between 15 December 2011 and 30 April 2012, both triggered by explosives and natural avalanches. The entire time period was labeled and split into an SVM training and a testing phase. The training phase consisted of 26 days between 15 December 2011 and 10 January 2012 during which 29 avalanches were observed. The rest of the winter period, from 11 January 2012 until the end of April 2012, was used as the testing period. Due to several days of system failure during this period, the testing phase consisted of 73 days with 30 observed avalanches.

For the training of the SVM, features from 20 out of the 29 avalanche events (700 positive feature vectors) were selected from the training period; these 20 events exhibited the strongest avalanche signal characteristics. Since more than 99.9% of all data are not assigned to avalanche events, subsampling of negative feature vectors was necessary to obtain the same number of feature vectors as for positive events. Instead of pure random subsampling from the entire training period, stratified subsampling was applied by manually selecting specific negative events (series of negatively classified feature vectors) which tend to resemble avalanches (similar signal length and properties) (Rubin et al., 2012). This leads to a better discrimination of positive feature vectors from negative ones during the training phase of the SVM and to a significant reduction of false positive classifications compared to pure random subsampling.

A linear SVM was trained yielding a 10-fold cross validated classification error rate of 8.6% (91.4% classification accuracy). This error rate applies to the case where all features of Table 1 were taken into account for the training. The right column of Table 1 further shows the cross-validation error rates if the classifier was trained by using only a single feature. The cross-validation error rate was used to optimize the parameters for the feature extraction (see Section 2.3). It further indicates the sensitivity of the classification with respect to a certain feature.

3.2. SVM test and avalanche detection accuracy

Avalanche detection accuracy was assessed during the testing period by counting the number of detected events, false detections and missed avalanches. Event observation was spatially restricted to the avalanche controlled area where complete avalanche occurrence records were available. Spatial restriction was achieved by applying an azimuth angle post filter to the classification output of the SVM, accepting only events with azimuth angles which intersect with the avalanche controlled area (azimuth angle range: −20° to −70°). Fig. 2 shows the evaluation of the classification results during the testing period. Fig. 2a displays the relative amounts of the true positive (detected avalanches) and false negative (missed avalanches) detections among all observed avalanches (30 events). 17 events (57%) were detected and 13 events (43%) were missed. This rather low detection rate can be explained by the relatively high number of small avalanches during the testing period. Small avalanches do not generate sufficient infrasound and thus it is often not possible to detect them in the far field. Especially in the spring season, the number of small (wet snow) avalanches is expected to be higher than in the winter season, leading to an increased amount of missed events.

Overall, the classifier detected 31 events in the avalanche controlled area during the testing period. Fig. 2b shows the relative amounts of true positive and false positive detections among all these detected events. True positive events were further split into “confirmed” and “likely” events. While “confirmed” refers to events which were in fact observed and recorded by the local observer, “likely” refers to events which were not reported but show clear avalanche properties upon manual inspection of the data.

The SVM classifier was compared to the default classifier provided by the manufacturer of the commercial infrasound system. This classifier is based on thresholds applied to specific signal features which have been empirically derived and optimized in-situ. In particular, the

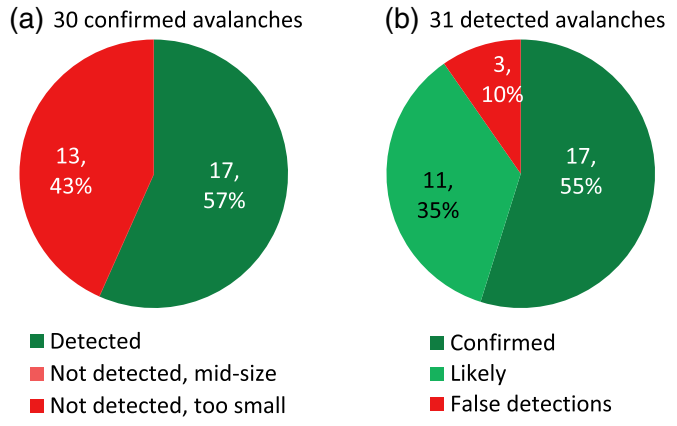

Fig. 2. Classification results of the test period for the proposed SVM-based classifier in the avalanche controlled area. (a) Out of 30 confirmed avalanches, 17 avalanches (57%) were detected and 13 (43%) were missed because they were too small. (b) The classifier detected 31 avalanches in total of which 17 (55%) are confirmed events and 11 (35%) are likely events. There were 3 (10%) false detections.

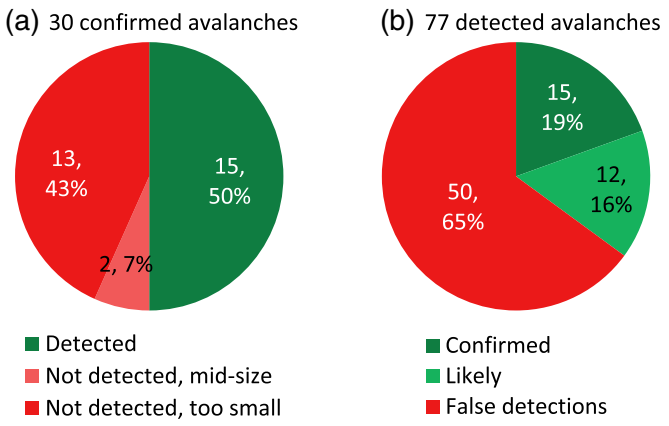

Fig. 3. Classification results of the test period for the threshold-based classifier in the avalanche controlled area. (a) Out of 30 confirmed avalanches, 15 avalanches (50%) were detected, 2 (7%) were not detected and 13 (43%) were not detected because they were too small. (b) The classifier detected totally 77 avalanches of which 15 (19%) are confirmed events and 12 (16%) are likely events. There were 50 (65%) false detections.

method classifies avalanche events based on the event duration time, mean azimuth and elevation angles (they must be in the range intersecting with the avalanche control site), extension (minimum and maximum) and sign of the azimuth angle, correlation coefficients of the sensor signals and infrasonic power in a low frequency band. Features are updated every 2 s (Chritin et al., 1996; IAV, 2006). Unfortunately, no complete description of the method used in this commercial product is available. Fig. 3 shows the classification results of the threshold-based classifier for the testing period. The false detection rate (Fig. 3b) is 65% compared to 10% for the SVM-based classifier (Fig. 2b). At the same time, two confirmed avalanches, which were detected by the SVM method, were missed by the default classifier. The significantly increased false alarm rate of the threshold-based classifier may suggest that the avalanche detection sensitivity with this method is simply higher than with the SVM-based classifier. However, the fact that the SVM-based classifier detects more confirmed avalanches than the threshold-based classifier and at the same time exhibits less false alarms clearly indicates that this is not the case.

3.3. Avalanche activity assessment

Avalanche activity was assessed in the entire area around the sensor array by removing the azimuth angle filter applied in the previous section. In total, 50 avalanches (i.e. 19 additional avalanches) were detected during the entire testing phase and over an azimuth angle range from −180° to 180°. Manual inspection of the data revealed that these additional 19 avalanches were in fact "likely" avalanche events. Fig. 4a shows the azimuth angle map, which qualitatively displays the directions of incidence of each recorded avalanche event. The plot shows that the majority of avalanches originated from the avalanche controlled area, which is located north-west of the sensor array (marked with red arrows). Fig. 4b further shows the avalanche activity in the entire area during the testing phase. Events are again divided into avalanches which originated from the avalanche controlled area (red) and from other directions. Avalanche events outside of the controlled area often occurred close in time to events in the controlled area, indicating periods of high avalanche activity.

4. Discussion and conclusions

Robust avalanche detection with an infrasonic sensor array using a support vector machine was demonstrated. In order to correctly assess detection accuracy, the proposed method was evaluated in an avalanche controlled area (distance range of 2–3 km). The method was compared to a threshold-based classifier provided by the manufacturer of the commercial system. The false alarm rate, which has been a major issue with this classifier, was significantly decreased from 50 (65%) to 3 (10%) events with the proposed method. At the same time, the detection rate slightly improved from 50% to 57%. The rather small detection rate of 57% (43% missed events) is due to the relatively large amount of small and most likely wet-snow avalanches during the testing phase. These avalanches only emitted small infrasonic power and since the sensor array was operating in the far-field, the sensitivity may not have been high enough to detect these events (Kogelnig et al., 2011).

SVM-based and threshold based classification are fundamentally different methods in terms of implementation and optimization. It would not be appropriate to conclude that SVM-based classifiers are in general superior to threshold-based classifiers. However in this study where the two classifiers were individually optimized to the specific system and site it was demonstrated that the SVM-based classifier clearly outperformed the threshold-based classifier. In general, threshold-based methods enable the definition of feature margins which can easily be controlled and optimized. On the other hand, mutual dependencies between features are difficult to establish correctly (Kotsiantis, 2007). By neglecting mutual dependencies, positive and negative events may not be separable and the classifier may either have a low detection or a high false detection rate. SVM-based classifiers automatically determine classification boundaries by a statistical analysis of training data and can take into account mutual dependencies between features. On

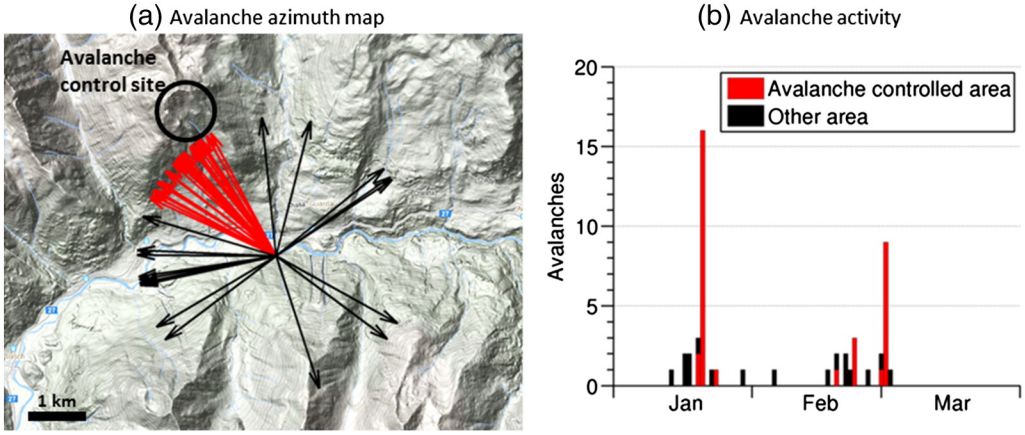

Fig. 4. Avalanche activity assessment during the testing period (11 January–30 April 2012). (a) The avalanche azimuth map shows the direction of origin of the detected avalanches. The center point corresponds to the position of the sensor array. Each arrow corresponds to a detected avalanche. Avalanches which originated from the avalanche control site are colored in red. (b) The avalanche activity plot shows the time of occurrence of the detected avalanche events. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the other hand, SVM methods do not offer direct access to classification margins and thus no easy control of the classification. The user must carefully select the training data containing relevant positive and negative events. In particular for avalanche detection, a certain amount of avalanches are required which should exhibit some variability in their signal properties to account for different characteristics (e.g., avalanche size or length). This may take half of a winter season or even more, depending on snow and meteorological conditions. However, setting up an infrasonic system always involves extensive adjustments to the environmental conditions (position, ambient noise sources, avalanche sizes, etc.) for any kind of classification method.

Feature extraction parameters have been derived by minimizing the error rate of the classification when using only a specific feature. While these parameters are likely to be highly dependent on general system properties such as system geometry, microphone input response, ambient noise or the number of sensors, the optimization procedure can be used on any system.

Despite the availability of comprehensive avalanche observation data in the avalanche controlled area, this study may still contain an unknown number of undetected avalanche events (e.g. small slides).

The proposed method was implemented and applied "off-line", i.e., by analyzing recorded data from the past. In a real-time environment, where the algorithm would run embedded on the system, the delay time from the avalanche release until its detection would mainly depend on the filter lengths. Taking into account all filters used for the proposed method, a minimum delay time of approx. 15 s would be realistic. A total delay time from avalanche release to a warning system (e.g., a message to the avalanche control staff, or the triggering of a traffic light) is estimated to be at least 20 s, which from an operational perspective may be acceptable, depending on the size and speed of the avalanche and the length of the avalanche path.

Infrasonic array-based avalanche detection is a promising technique if high detection reliability can be guaranteed. While sensor performance and sensitivity has been greatly improved in the past, the bottleneck of today's integrated detection systems remains in the signal processing part, in particular in event classification. This study presented a robust method based on a machine learning technique to reduce false alarms and improve overall detection accuracy.

Acknowledgments

We would like to thank Lorenz Meier for fruitful discussions about infrasonic sensor-array systems, Peder Caviezel for meticulously recording the avalanche events and Vincent Critin for providing technical specifications about the system.

References

Adam, V., Chritin, V., Rossi, M., van Lancker, E., 1998. Infrasonic monitoring of snow-avalanche activity: what do we know and where do we go from here? Ann. Glaciol. 26, 324–328.
Bedard, A., 1989. Detection of avalanches using atmospheric infrasound. Proceedings of the 57th Annual Western Snow Conference. Fort Collins CO, USA, pp. 52–58.
Bessason, B., Eiríksson, G., Thórarinsson, Ó., Thórarinsson, A., Einarsson, S., 2007. Automatic detection of avalanches and debris flows by seismic methods. J. Glaciol. 53, 461–472.
Burges, C., 1998. A tutorial on support vector machines for pattern recognition. Data Min. Knowl. Discov. 2, 121–167.
Cannata, A., Montalto, P., Aliotta, M., Cassisi, C., Pulvirenti, A., Privitera, E., Patanè, D., 2011. Clustering and classification of infrasonic events at Mount Etna using pattern recognition techniques. Geophys. J. Int. 185, 253–264. http://dx.doi.org/10.1111/j.1365-246X.2011.04951.x.
Christie, D., Vivas Veloso, J., Campus, P., Bell, M., Hoffmann, T., Langlois, A., Martysevich, P., Demirovic, E., Carvalho, J., 2001. Detection of atmospheric nuclear explosions: the infrasound component of the International Monitoring System. Kerntechnik 66, 96–101.
Chritin, V., Adam, V., Rossi, M., Bolognesi, R., 1996. Acoustic detection system for operational avalanche forecasting. Acta Acoust. 82, 173.
Cristianini, N., Shawe-Taylor, J., 2000. An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods. Cambridge University Press.
Gauer, P., Kern, M., Kristensen, K., Lied, K., Rammer, L., Schreiber, H., 2007. On pulsed Doppler radar measurements of avalanches and their implication to avalanche dynamics. Cold Reg. Sci. Technol. 50, 55–71. http://dx.doi.org/10.1016/j.coldregions.2007.03.009.
IAV, 2006. Arfang: Infrasonic Avalanche Monitoring.
Johnson, D., Dudgeon, D., 1992. Array signal processing: concepts and techniques. 1st ed. Prentice Hall (February 11, 1993).
Knapp, C., Carter, G., 1976. The generalized correlation method for estimation of time delay. IEEE Trans. Acoust. Speech Signal Process. ASSP-24, 320–327.
Kogelnig, A., Suriñach, E., Vilajosana, I., Hübl, J., Sovilla, B., Hiller, M., Dufour, F., 2011. On the complementariness of infrasound and seismic sensors for monitoring snow avalanches. Nat. Hazards Earth Syst. Sci. 11, 2355–2370. http://dx.doi.org/10.5194/nhess-11-2355-2011.
Kogelnig, A., Wyssen, S., Pichler, J., 2012. Artificial release and detection of avalanches: managing avalanche risk on traffic infrastructure, a case study from Austria. International Snow Science Workshop. Anchorage AK, USA, pp. 535–540.
Kotsiantis, S., 2007. Supervised machine learning: a review of classification techniques. Informatica 31, 249–268.
Kotsiantis, S., Zaharakis, I., Pintelas, P., 2007. Machine learning: a review of classification and combining techniques. Artif. Intell. Rev. 26, 159–190. http://dx.doi.org/10.1007/s10462-007-9052-3.
Lacroix, P., Grasso, J.-R., Roulle, J., Giraud, G., Goetz, D., Morin, S., Helmstetter, A., 2012. Monitoring of snow avalanches using a seismic array: location, speed estimation, and relationships to meteorological variables. J. Geophys. Res. 117, F01034. http://dx.doi.org/10.1029/2011JF002106.
Larsen, S., Salberg, A., Solberg, R., 2010. Automatic Detection of Avalanches in High-Resolution Optical Satellite Data — Results From the ESA avalRS Project's Feasibility Study on Automated Avalanche Detection. Oslo, Norway.
Lato, M., Frauenfelder, R., Bühler, Y., 2012. Automated detection of snow avalanche deposits: segmentation and classification of optical remote sensing imagery. Nat. Hazards Earth Syst. Sci. 12, 2893–2906. http://dx.doi.org/10.5194/nhess-12-1-2012.
Le Pichon, A., Blanch, E., Hauchecorne, A., 2009. Infrasound Monitoring for Atmospheric Studies. Springer Netherlands, Dordrecht, Netherlands http://dx.doi.org/10.1007/978-1-4020-9508-5.
Leprettre, B., Navarre, J., Taillefer, A., 1996. First results from a pre-operational system for automatic detection and recognition of seismic signals associated with avalanches. J. Glaciol. 42, 352–363.
McClung, D., Schaerer, P., 2006. The Avalanche Handbook. The Mountaineers Books, Seattle WA, USA.
Meier, L., Lussi, D., 2010. Remote detection of snow avalanches in Switzerland using infrasound, Doppler radars and geophones. International Snow Science Workshop. Lake Tahoe CA, USA, pp. 7–12.
Ripepe, M., Marchetti, E., Ulivieri, G., 2007. Infrasonic monitoring at Stromboli volcano during the 2003 effusive eruption: insights on the explosive and degassing process of an open conduit system. J. Geophys. Res. Solid Earth 112, B09207. http://dx.doi.org/10.1029/2006JB004613.
Rubin, M., Camp, T., van Herwijnen, A., Schweizer, J., 2012. Automatically detecting avalanche events in passive seismic data. International Conference on Machine Learning and Applications. IEEE, pp. 13–20 http://dx.doi.org/10.1109/ICMLA.2012.12.
Salm, B., Gubler, H., 1985. Measurement and analysis of the motion of dense flow avalanches. Ann. Glaciol. 6, 26–34.
Schaerer, P., Salway, A., 1980. Seismic and impact-pressure monitoring of flowing avalanches. J. Glaciol. 26, 179–187.
Schimmel, A., Hübl, J., 2013. Automatic detection of avalanches using infrasound and seismic signals. In: Naaim-Bouvet, F., Durand, Y., Lambert, R. (Eds.), International Snow Science Workshop. Grenoble, France, pp. 904–908.
Schweizer, J., Jamieson, J., 2010. Snowpack tests for assessing snow-slope instability. Ann. Glaciol. 51, 187–194.
Schweizer, J., van Herwijnen, A., 2013. Can near real-time avalanche occurrence data improve avalanche forecasting? In: Naaim-Bouvet, F., Durand, Y., Lambert, R. (Eds.), International Snow Science Workshop. Grenoble, France, pp. 195–198.
Scott, E., Hayward, C., Colgan, T., Hamann, J., Kubichek, R., Pierre, J., Yount, J., 2006. Practical implementation of avalanche infrasound monitoring technology for operational utilization near Teton Pass Wyoming. In: Gleason, A. (Ed.), International Snow Science Workshop. Telluride CO, USA, pp. 714–723.
Scott, E., Hayward, C., Kubichek, R., Hamann, J., Pierre, J., Comey, B., Mendenhall, T., 2007. Single and multiple sensor identification of avalanche-generated infrasound. Cold Reg. Sci. Technol. 47, 159–170. http://dx.doi.org/10.1016/j.coldregions.2006.08.005.
Suriñach, E., Sabot, F., Furdada, G., Vilaplana, J., 2000. Study of seismic signals of artificially released snow avalanches for monitoring purposes. Phys. Chem. Earth Part B 25, 721–727.
Ulivieri, G., Marchetti, E., Ripepe, M., Chiambretti, I., De Rosa, G., Segor, V., 2011. Monitoring snow avalanches in Northwestern Italian Alps using an infrasound array. Cold Reg. Sci. Technol. 69, 177–183. http://dx.doi.org/10.1016/j.coldregions.2011.09.006.
Van Herwijnen, A., Schweizer, J., 2011a. Seismic sensor array for monitoring an avalanche start zone: design, deployment and preliminary results. J. Glaciol. 57, 267–276. http://dx.doi.org/10.3189/002214311796405933.
Van Herwijnen, A., Schweizer, J., 2011b. Monitoring avalanche activity using a seismic sensor. Cold Reg. Sci. Technol. 69, 165–176. http://dx.doi.org/10.1016/j.coldregions.2011.06.008.
Van Lancker, E., 2001. Acoustic goniometry: a spatio-temporal approach. PhD thesis. Ecole Polytechnique Fédéral de Lausanne (EPFL), Lausanne, Switzerland.
Vapnik, V., 2000. The nature of statistical learning theory. Data Mining and Knowledge Discovery 2nd ed. Springer Verlag.
Vriend, N., McElwaine, J., Sovilla, B., Keylock, C., Ash, M., Brennan, P., 2013. High-resolution radar measurements of snow avalanches. Geophys. Res. Lett. 40, 727–731. http://dx.doi.org/10.1002/grl.50134.