Automatic detection of avalanches using a combined array classification and localization

Matthias Heck[1], Alec van Herwijnen[1], Conny Hammer[2], Manuel Hobiger[2], Jürg Schweizer[1], and Donat Fäh[2]

[1]WSL Institute for Snow and Avalanche Research SLF, Davos
[2]Swiss Seismological Service SED, ETH Zurich, Zurich

Correspondence: Matthias Heck (matthias.heck@slf.ch)

Abstract.

We used continuous data from a seismic monitoring system to automatically determine the avalanche activity at a remote field site above Davos, Switzerland. The approach is based on combining a machine learning algorithm with array processing techniques to provide an operational method capable of near real-time classification. First, we used a recently developed method based on hidden Markos models (HMMs) to automatically identify events in continuous seismic data by using only a single training event. For the 2016-2017 winter period, this resulted in 117 events. Second, to eliminate falsely classified events such as airplanes and local earthquakes, we implemented an additional HMM based classifier at a second array $14\,\mathrm{km}$ away. By cross-checking the results of both arrays, we reduced the number of false classifications by about $50\,\%$. In a third and final step we used multiple signal classifications (MUSIC), an array processing technique, to determine the direction of the source. Since snow avalanches recorded at our arrays typically generate signals with small changes in source direction, events with large changes were dismissed as false detections. From the 117 initially detected events during the 4-month period, our classification workflow removed 96 events as false classifications. The majority of the remaining 21 events were on 9 and 10 March 2017, in line with visual avalanche observations in the region of Davos. Our results suggest that the presented classification workflow could be used to identify major avalanche periods and highlight the importance of array processing techniques for the automatic classification of avalanches in seismic data.

1 Introduction

During the winter seasons, snow avalanches are a common threat in mountainous regions. Avalanche warning services therefore inform the public of the avalanche danger. To assess the danger, warning services rely on information about the snowpack, amount of new snow, weather conditions and avalanche activity (e.g. McClung and Schaerer, 2006). Whereas the first three parameters can be measured or modeled, avalanche activity data are often hard to obtain, especially during snow storms or at night. Monitoring systems can possibly fill this gap by providing information on avalanche activity independent of the time of day or visibility.

Snow avalanches, like any other mass movement, generate seismic and infrasound waves (e.g. van Herwijnen and Schweizer, 2011b; Suriñach et al., 2005; Marchetti et al., 2015). Seismic signals of avalanches show some common characteristics, including a spindle shaped envelope of the time series (Nishimura and Izumi, 1997) and a typical frequency content between 2 and 30 Hz (Schaerer and Salway, 1980; Suriñach et al., 2001). Several classification approaches were therefore developed to automatically detect avalanches in seismic data. Leprettre et al. (1996) used a fuzzy logic approach to distinguish between different types of signals. Bessason et al. (2007) applied a nearest neighbor approach and detected 65% of all confirmed avalanches. Rubin et al. (2012) compared 12 machine learning algorithms, 10 of which were able to detect at least 90% of all manually identified avalanches. While these machine learning methods perform reasonably well in terms of detecting confirmed avalanche events, a large training data set is typically required and false alarm rates are generally rather high (Rubin et al., 2012).

An alternative machine learning approach was recently presented by Hammer et al. (2017). They used hidden Markov models (HMMs) to automatically detect large avalanches released in February 1999 in seismic data recorded by a single broadband station maintained by the Swiss Seismological Service (SED). HMMs use a sequence of multivariate Gaussian probability distributions to model observations (e.g. seismic time series). To determine the characteristics of the distributions (i.e. mean and covariance), classical HMMs also require a large number of training sets for each event class (e.g. avalanche, airplane or earthquake). This classical approach was successfully used to automatic identify seismic events in continuous seismic data (Ohrnberger, 2001; Beyreuther et al., 2012). Avalanches, however, are relatively rare events and obtaining a large set of training events is time consuming. To circumvent this, Hammer et al. (2012) developed an approach exploiting the abundance of data containing mainly background signals to obtain general wave-field properties. From these properties, a widespread background model was learned and only one training event was required. In contrast to the classical HMM approach, the classification system thus consists of a background model and one event model for each implemented event class. Using this approach, Hammer et al. (2017) were able to identify 43 destructive avalanches during an exceptional 5-day avalanche period in February 1999 within a radius of $30\,\mathrm{km}$ of the broadband seismic station. Heck et al. (2018a) recently adapted this HMM approach to automatically detect smaller avalanches in data recorded by a seismic array consisting of seven less sensitive vertical geophones during the winter season 2009-2010. Despite the large differences in model performance for the individual sensors, their model performed best when pooling the data from the entire array and requiring a minimal event duration for the detections. These results highlighted that array based processing is likely to improve the automatic detection of avalanches in continuous seismic data.

Array processing techniques have been used to exploit infrasound signals for the automatic detection of avalanches (Scott et al., 2007; Marchetti et al., 2015; Thüring et al., 2015). By comparing the back-azimuth with the directions of known avalanche paths, possible avalanche events were identified (Marchetti et al., 2015). Thüring et al. (2015) used a similar approach for the automatic detection, but relied on support vector machines (SVM), a machine learning algorithm. The success of array processing techniques on infrasound signals led to the development of operational avalanche detection systems to automatically identify avalanches (Steinkogler et al., 2016). A recent comprehensive study on the performance of these systems has shown that in the absence of major topographic barriers, infrasound avalanche detection systems relying on array processing techniques are well suited to reliably monitor larger avalanches up to a distance of 3 to 4 kilometers (Mayer et al., 2018).

Array processing techniques have also been used to locate the source of avalanches in seismic data. Lacroix et al. (2012) implemented a beam-forming approach and were able to assign recorded avalanches to three known avalanche paths. Heck et al. (2018b) compared a beam-forming method with a multiple signal classification (MUSIC) approach (Schmidt, 1986) and obtained better results with the latter. The MUSIC method is based on the covariance matrix of all sensors, whereas beam-forming methods rely on pair-wise cross-correlation (Schmidt, 1986; Rost and Thomas, 2002). Heck et al. (2018b) subsequently applied this method to manually identified avalanches during a two-day period in March 2017 and were able to reconstruct the avalanche path of several recorded events. They concluded that their seismic array mostly recorded infrasound signals due to the limited distance between the sensors. While both Lacroix et al. (2012) and Heck et al. (2018b) showed that avalanches within a distance of approximately $3\,\mathrm{km}$ of their seismic monitoring systems were detected, seismic array processing techniques have not yet been used for the automatic detection of avalanches in seismic data.

Our aim is to design a workflow to automatically identify avalanches in continuous seismic data. Our method consists of using the machine learning techniques based on hidden Markov models presented by Heck et al. (2018a) in combination with seismic array processing techniques to locate the source of avalanches presented by Heck et al. (2018b). The goal is to assess the performance of a fully automatic classifier in view of possible future operational use. To develop and test the automatic classification method, we used continuous seismic data recorded during the winter season 2016-2017 at two field sites above Davos, Switzerland.

2 Field site and instrumentation

During the 2016-2017 winter season, we installed two seismic arrays above Davos, Switzerland (Figure 1). The arrays were similar to the system described by van Herwijnen and Schweizer (2011a). The first array was deployed at the Dischma field site (yellow square in Figure 1), $14\,\mathrm{km}$ away from Davos at the end of a tributary valley (Heck et al., 2018b). The field site is a flat meadow at an elevation of $2000\,\mathrm{m}$ a.s.l. surrounded by mountain peaks which rise up to $3000\,\mathrm{m}$. The second array was deployed at the Wannengrat field site above Davos at $2400\,\mathrm{m}$ a.s.l. (red square in Figure 1). This field site is surrounded by several avalanche starting zones (van Herwijnen and Schweizer, 2011a). Both arrays consisted of a $300\,\mathrm{m}$ long string with 7 vertical component geophones with an natural frequency of $4.5\,\mathrm{Hz}$. The sensors of the Dischma array were buried about $50\,\mathrm{cm}$ into the ground whereas the sensors at the Wannengrat field site were attached to rocks using an anchor. For each array, the sensors were circularly arranged (Figure 1). The maximum distance between two sensors at the Dischma and Wannengrat field site was $64\,\mathrm{m}$ and $74\,\mathrm{m}$, respectively, and the average distance was $36\,\mathrm{m}$ and $45\,\mathrm{m}$, respectively.

The instrumentation and data logging systems were identical for both arrays. Data were continuously recorded at a sampling rate of $500\,\mathrm{Hz}$. Due to technical problems, only two sensors of the Wannengrat array recorded data throughout the entire winter (4 and 5 in Figure 1b) and no data were collected between 12 and 20 January 2017. Both field sites were also equipped with several automatic cameras (8 at Dischma, 6 at Wannengrat) to monitor the surrounding slopes. Images were recorded every 10 minutes throughout the winter (Figure 2). In addition, we also performed a field survey at the Dischma site on 15 March 2017 shortly after a period of high avalanche activity to identify avalanches and map their outlines (Figure 3) (Heck et al., 2018b).
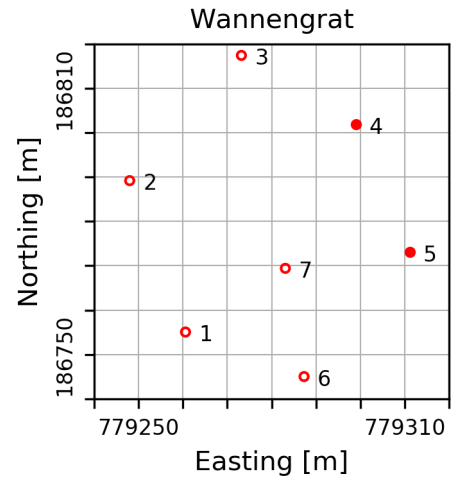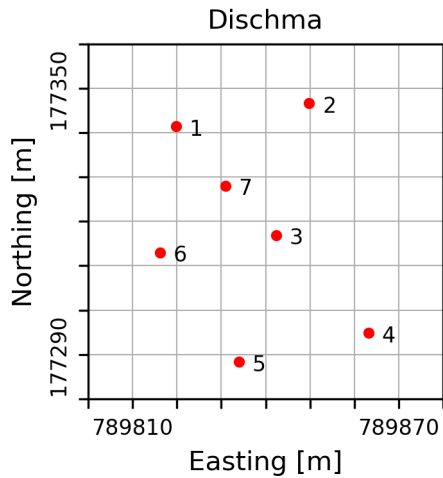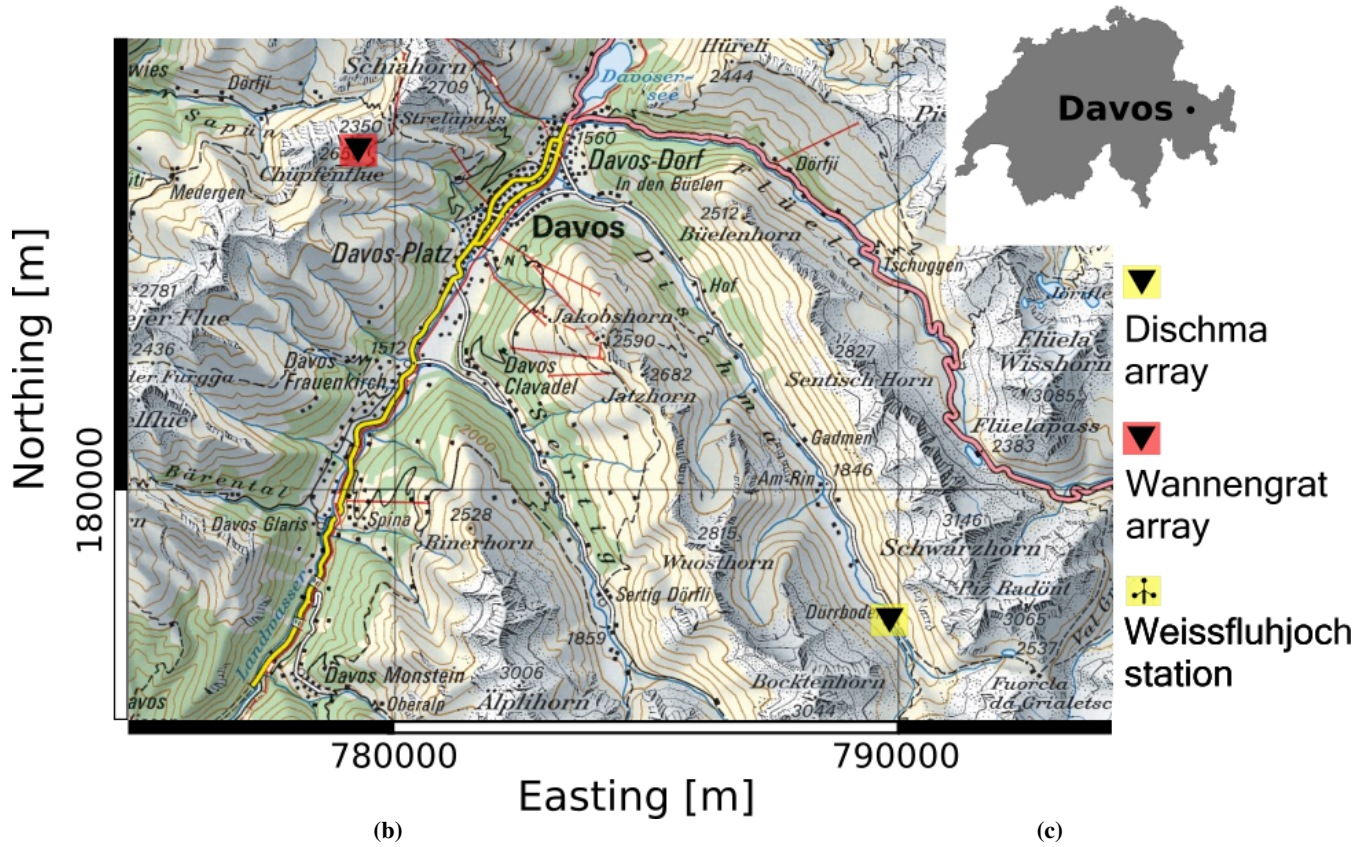
3

Figure 1. a) Map of the area of Davos, Switzerland. The two seismic arrays are indicated by a black triangle on colored background. Red represents the Wannengrat array, yellow the Dischma array. The wind wheel indicates the location of the Weissfluhjoch weather station. (b) Deployment geometry of the Dischma array. (c) Deployment geometry of the Wannengrat array. The open red circles indicate positions of malfunctioning sensors during the winter 2017. Reproduced by permission of swisstopo (JA100118).
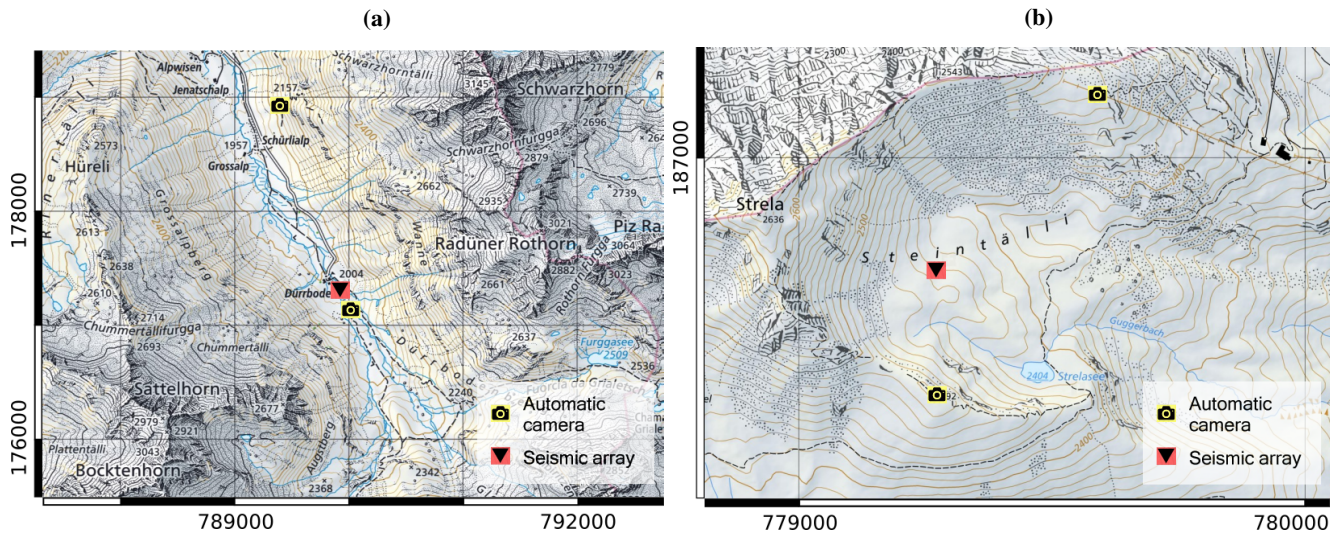
Figure 2. (a) Detailed map of the Dischma field site showing the position of the seismic array and automatic cameras. (b) Same for the Wannengrat field site.

3 Methods

The overall goal was to develop a processing workflow to automatically identify avalanches in seismic data from the Dischma field site by combining HMMs with seismic array processing techniques. The developed workflow consists of five steps which are described in more detail below (Figure 4): (1) pre-processing, (2) feature calculation, (3) HMM construction, (4) classification and (5) post-processing.

3.1 Data pre-processing

The continuous seismic data mostly consisted of noise, which for the current application was of little interest. We therefore applied a simple threshold based event detector to reduce the total amount of data (Heck et al., 2018a). For a window i with a length of 1024 samples, we determined the mean absolute amplitude A_i. When $A_i \geq 5\overline{A}$, with $\overline{A}$ the daily mean amplitude, the data within the window were kept. If the amplitude threshold for the following window was also reached, data were concatenated. Furthermore, a section of $t = 60\,\mathrm{s}$ before and after the threshold passing were kept to ensure that the onset and coda of events were incorporated. Doing so, data were reduced by $80\,\%$ to data windows of various lengths. In addition, we filtered the data using a 4th order Butterworth bandpass filter with corner frequencies of 1 and $50\,\mathrm{Hz}$.

3.2 Feature calculation

Raw seismic time series are not suited for the HMM classification, as information which characterizes seismic signals generated by avalanches in the time and frequency domain cannot be exploited. We therefore used specific features of the seismic time series as input for the HMMs. Features represent different aspects of the time series such as spectral, temporal or polarization

5

Figure 3. Picture of the field site of Dischma facing to the south. It was taken on 15 March 2017 shortly after a period of high avalanche activity and several recent avalanches were observed (red areas). The triangle indicates the location of the seismic array and the camera icon the location of the automatic camera.

characteristics. These are calculated using a sliding window and therefore the time series is represented in a compressed form. Since we used data from single component geophones, we only used the spectral and temporal features suggested by Heck et al. (2018a):

- – Central frequency (Barnes, 1993)

5
- – Dominant frequency (Kramer, 1996)

- – Instantaneous bandwidth (Barnes, 1993)

- – Instantaneous frequency (Taner et al., 1979)

- – Cepstral coefficients (Kanasewich, 1981)
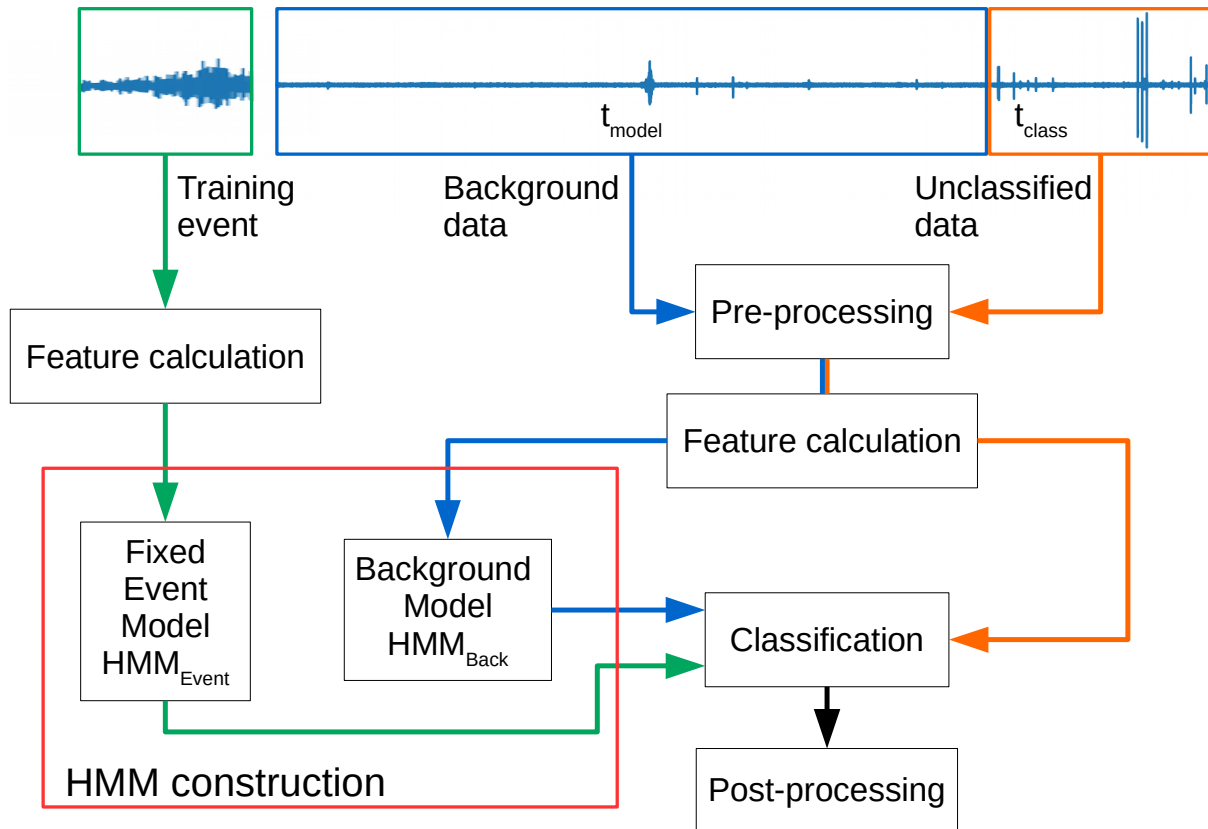
- – Half-octave bands (Joswig, 1994)

Figure 4. Flow chart of the classification process. Green lines show the construction of the event model. Blue lines show the construction of the background model. The orange lines show, how the data to be classified are processed.

To calculate the features from the pre-processed data, we used a sliding window of width $w = 512$ samples and a step size of 25 samples, resulting in an overlap of 97%. In total, we used 6 half-octave bands for the classification and the first half-octave band had a central frequency of $3.9\,\text{Hz}$.

3.3 HMM construction

5 The hidden Markov models we used to automatically identify avalanches in the continuous seismic consisted of a background model and an event model for each event class (Hammer et al., 2017; Heck et al., 2018a). In our application we used only one event class, either avalanches at the Dischma field site or airplanes at the Wannengrat field site.

The background model was learned by using features extracted from the pre-processed data. Previous work has shown that a background model which is regularly updated provides better classification results throughout a season (e.g. Heck et al., 2018a).

10 In our case, we therefore used the pre-processed data from a $24\,\text{h}$ window to train the background model and recalculated it

every hour, i.e. we used a sliding window of length $t_{\mathrm{model}} = 24\,\mathrm{h}$ and then shifted it forward by one hour (background data in Figure 4).

In contrast to the background model, the event model was only calculated once for the entire season (training event in Figure 4). For the Dischma field site, our training event consisted of a signal generated by an avalanche that had released on 9 March 5 2017 at 06:47 (Figure 5a), an unambiguous event that was analyzed in detail by Heck et al. (2018b). For the Wannengrat field site, our training event consisted of a signal generated by an airplane on 31 January 2017 at 21:46 (Figure 5b).
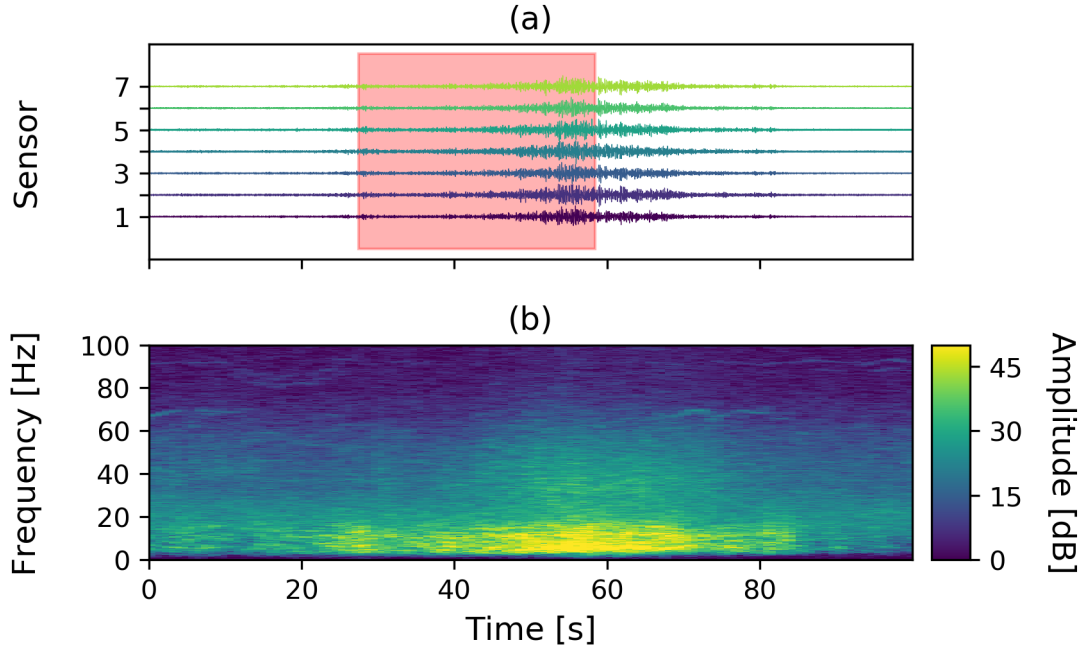

Figure 5. Avalanche released on 9 March 2017 at 06:47 used as training event for the classifier at the Dischma field site. a) time series for the 7 sensors. The red area indicates the part of the time series used as training event. b) corresponding spectrogram of the seismic time series.

3.4 Classification

To classify unknown data, we used a window of length $t_{\mathrm{class}} = 1\,\mathrm{h}$ immediately following the $24\,\mathrm{h}$ long training window (Unclassified data in Figure 4). These data were pre-processed and features were calculated. Based on these features, the 10 HMM classification process then calculated the likelihood that an unknown data stream was generated by a specific event class (Hammer et al., 2012, 2013). The classification was performed for each individual sensor.

3.5 Post-processing

Since the classification algorithm is not perfect, several post-processing steps were applied to reduce the number of false detections. Detections with a duration $\leq 12\,\mathrm{s}$ were dismissed (Heck et al., 2018a), and we only considered an event as an

avalanche when it was detected by at least five sensors of the array, as proposed by Heck et al. (2018a). Two additional steps were also implemented in the classification workflow and are described in the following.

3.5.1 Combined array detection

Avalanche signals are typically only detected within a radius of $3\,\mathrm{km}$ of our seismic arrays (van Herwijnen and Schweizer, 2011b; Heck et al., 2018b). Since the distance between the Dischma and the Wannengrat array was $14\,\mathrm{km}$, it is therefore very unlikely that an avalanche was recorded at both arrays simultaneously. Signals that were recorded simultaneously at both arrays were thus most likely false detections. To remove such signals from the events automatically classified at the Dischma array, we implemented a second HMM trained with the airplane event from 31 January at 21:46 (Figure 5b) to classify the data recorded at the Wannengrat and subsequently using the same post-processing steps mentioned earlier. However, since only two sensors recorded data at the Wannengrat array, we only considered events detected by these two sensors. Classification results from the Dischma and Wannengrat arrays were then combined to remove all events recorded simultaneously at both arrays, which means that the time difference between the detections at both field sites was less than $1\,\mathrm{min}$.

3.5.2 Source localization

Array processing methods provide information on the signal source and can provide additional parameters to classify events as avalanches, as is commonly done when exploiting infrasound signals (Scott et al., 2007; Marchetti et al., 2015; Thüring et al., 2015). To locate the source of events, we used the MUSIC algorithm, as suggested by Heck et al. (2018b). The MUSIC method was applied to non-overlapping data windows for four frequency bands between 6 and $7.5\,\mathrm{Hz}$ to determine the back-azimuth angle and the apparent velocity of the incoming wave-field with time. The length of the windows was set to five cycles of the lower corner of the analysed frequency band, meaning that the data were divided into more windows for higher frequencies. By combining the back-azimuth values for all four frequency bands, we then applied a median smoothing filter on a moving window of $8\,\mathrm{s}$ to determine a median back-azimuth path with time, as in Heck et al. (2018b) (Figure 6a).

To decide whether a classification was associated with an avalanche or not, we applied a threshold value to the derivative of the median back-azimuth path. The assumption was that avalanche events have a relatively smooth median back-azimuth path with little variations, whereas false detections show larger variations in back-azimuth, which is also the case for earthquakes and airplanes for our specific array geometry. Indeed, the training event at the Dischma array had a duration of around $50\,\mathrm{s}$ and a median back-azimuth path with slight changes in the angle (black line in Figure 6 a) whereas before and after the event, the back-azimuths were more erratic, as expected for noise. For the training event, the derivative of the median back-azimuth path stayed below $10°$ for $50\,\mathrm{s}$, while before and after the event it was much larger (Figure 6b). Other avalanche events had very similar results (not shown). Events with derivatives of the median back-azimuth path smaller than $10°$ for a minimum duration of $20\,\mathrm{s}$ were then classified as avalanches. Since we used $8\,\mathrm{s}$ windows for the calculation of the median back-azimuth path, we had to increase the minimal event duration to $20\,\mathrm{s}$.
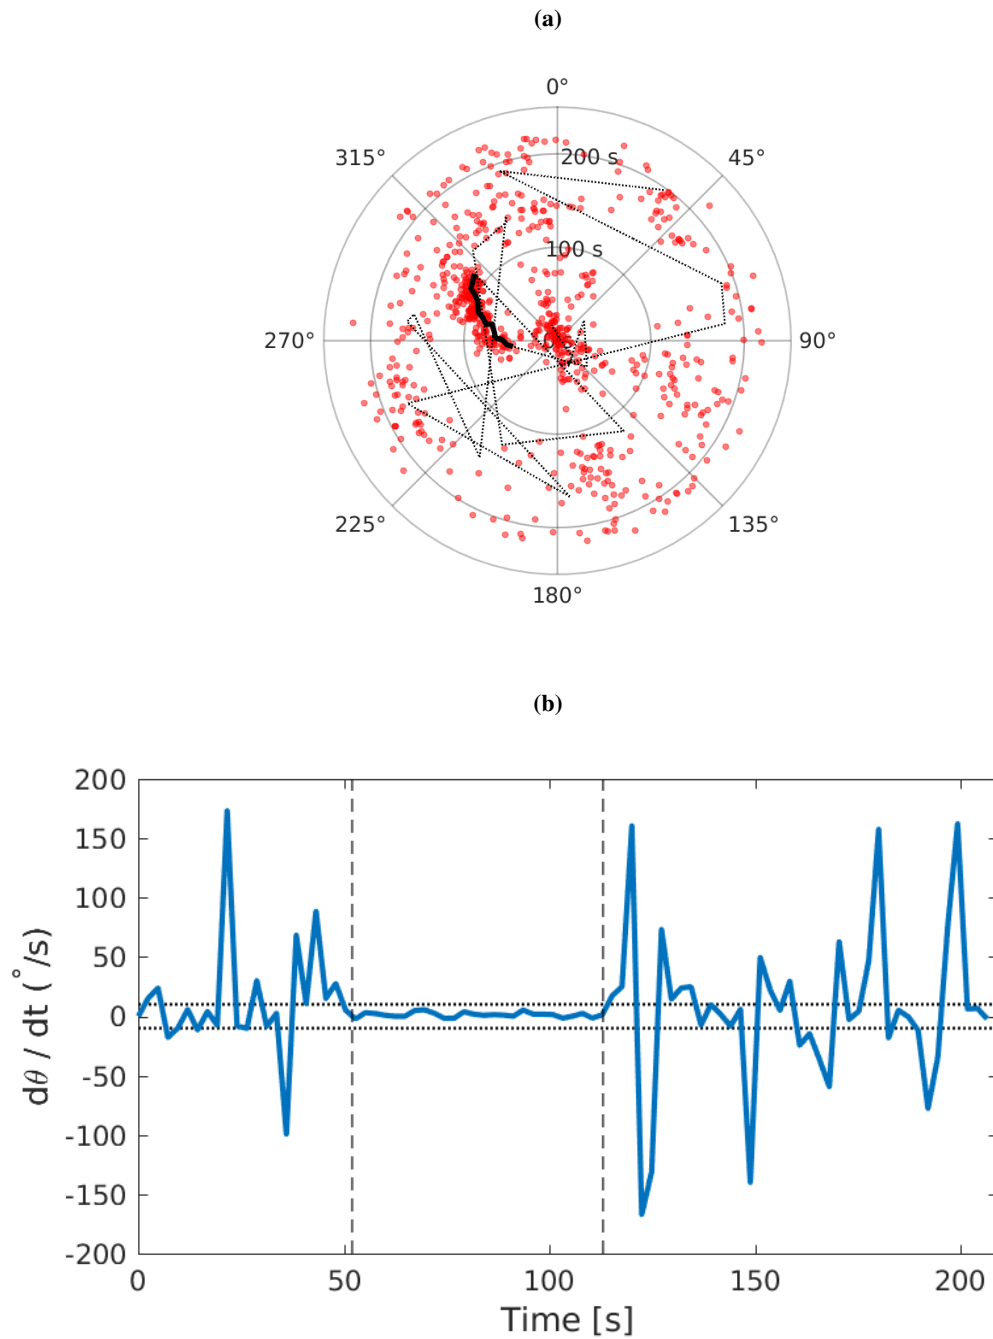
(a)

(b)

Figure 6. Localization results for an avalanche event recorded on 9 March 2017 at 6:47. a) polar plot representation of the back-azimuth calculated using the MUSIC method. Red dots are the back-azimuth values for a single time window. The black line represents the median back-azimuth path. The solid part of the line has variations below the threshold value for the derivative, whereas the dotted line refers to strong variations. b) derivative of median back-azimuth path. The dotted lines represents the threshold value of $10°$. The part between $52\,\mathrm{s}$ and $113\,\mathrm{s}$ corresponds to the solid line in a).

4 Results

We applied our automatic avalanche detection workflow on data recorded at the Dischma field site from 1 January to 30 April 2017. We then compared the results with an avalanche catalogue obtained by visual observations from local observers compiled by the avalanche warning service at the SLF. The visual observations in the area of Davos can be incomplete and cover an area much larger than the area monitored with our seismic system at the Dischma site. Therefore, comparison with this avalanche catalogue remains indicative.

4.1 Overview of the winter season

The winter period of 2016-2017 was relatively short and characterized by a below-average snow depth. First snowfalls were quite late in the season, in the middle of December, followed by four weeks without substantial precipitation and low temperatures. As the thin snowpack was subjected to large temperature gradients for an extended period of time, a poorly bonded layer of depth hoar formed at the base of the snow cover.

During the winter season, four pronounced snowfall periods occurred: between 1 and 15 January, around 1 February, from 1 to 10 March and around 15 April (increase of blue line in Figure 7). Each snowfall was associated with increased avalanche activity in the region of Davos, except the snowfall in April (red bars in Figure 7b).

In addition to these visual avalanche observations by local observers, we inspected pictures taken by the automatic cameras installed at the Dischma field site. Surprisingly, avalanche activity was very low in January and February, and only a few avalanches were identified on 1 February 2017. During the early March snow storm, visibility was poor and very few avalanches were identified on the images of the automatic cameras. However, once the storm had passed, the intensity of the avalanche cycle became clear as many avalanche deposits were visible on the images. Five days after the storm we mapped 24 avalanches within a $4\,\mathrm{km}$ radius of the Dischma field site (for more details see Heck et al., 2018b). Later in the season, we did not observe any more avalanches on the images. Thus, the avalanche cycle between 9 and 10 March was the most prominent avalanche period at the Dischma field site.

4.2 Classification performed at single array

Using the classifier trained with the avalanche event from 9 March 2017 at 06:47 (Figure 5a), we classified the data from each single sensor of the Dischma array and post-processed the results to remove events with a duration $\leq 12\,\mathrm{s}$ and all classifications that were classified by less than 5 sensors. This resulted in a total of 117 automatically detected events between 1 January and 30 April 2017 (Figure 8). A quarter of the events were detected by 5 sensors, a quarter by 6 and about half by 7 sensors.

Most events occurred on 9 and 10 March 2017. In addition, peaks in the middle of January and early April were visible. The peak in January corresponded with the avalanche activity period visually recorded in the region of Davos (Figure 7). For the peak in April, however, no avalanches were observed in the area of Davos. Furthermore, several single detections were distributed over the season showing no clear link to the avalanche observations in the region of Davos. We therefore visually
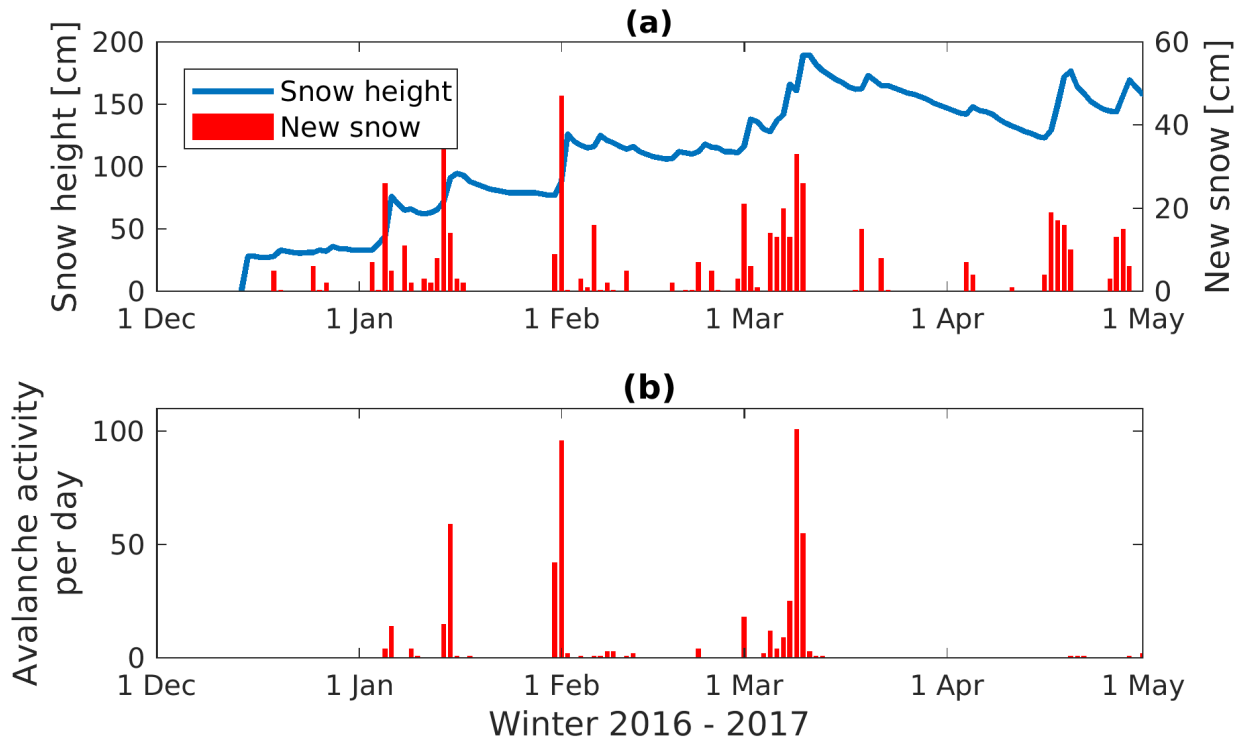
Figure 7. a) Snow height measured at the automatic weather station at Weissfluhjoch $12\,\mathrm{km}$ to the northwest of the Dischma array at $2540\,\mathrm{m}$ a.s.l. for the winter season 2016-2017. Orange bars are the height of new snow measured each day at 8:00 am. b) Number of avalanches observed per day in the region of Davos ($\sim 175\,\mathrm{km}^2$).

inspected the time series and spectrograms of each of the 117 classifications and found that the HMM also classified ~ 50 airplanes (Figure 9 a)) and regional earthquakes (Figure 9 b) as avalanches.

Although these misclassified events can be distinguished from avalanches through visual inspection (e.g. the sharp onset visible for earthquakes), the classifier identified these events as belonging to the avalanche class. We attribute these false classifications to the fact that these earthquake and airplane signals were more similar to the avalanche training event than to the background model. Indeed, the temporal trends in the features exhibited many similarities (Figure 10). Using different training events did not substantially change these results (not shown).

4.3 Classification performed at both arrays

The vast majority of the misclassifications were produced by two types of events: airplanes and earthquakes. Davos lies within an approaching corridor of the international airport Zürich and numerous commercial airplanes pass by every hour. Similar to avalanches, airplanes also have a moving source character. However, due to the high altitude (typically > 5 km) and fast movement of the source, airplanes are likely recorded at both arrays. Furthermore, based on the earthquake catalog of the
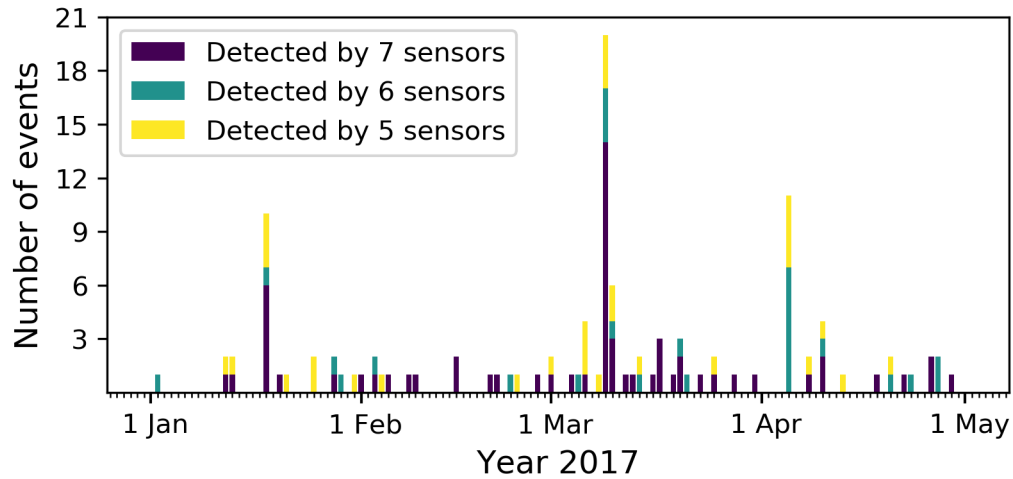
Figure 8. Classification results after post-processing (including voting-based classification) at the Dischma array. The colored bars indicate the number of classified events per day depending on the number of sensors: Violet bar indicates detections by 7 sensor, turquoise bars by 6 sensors and yellow bars by 5 sensors.
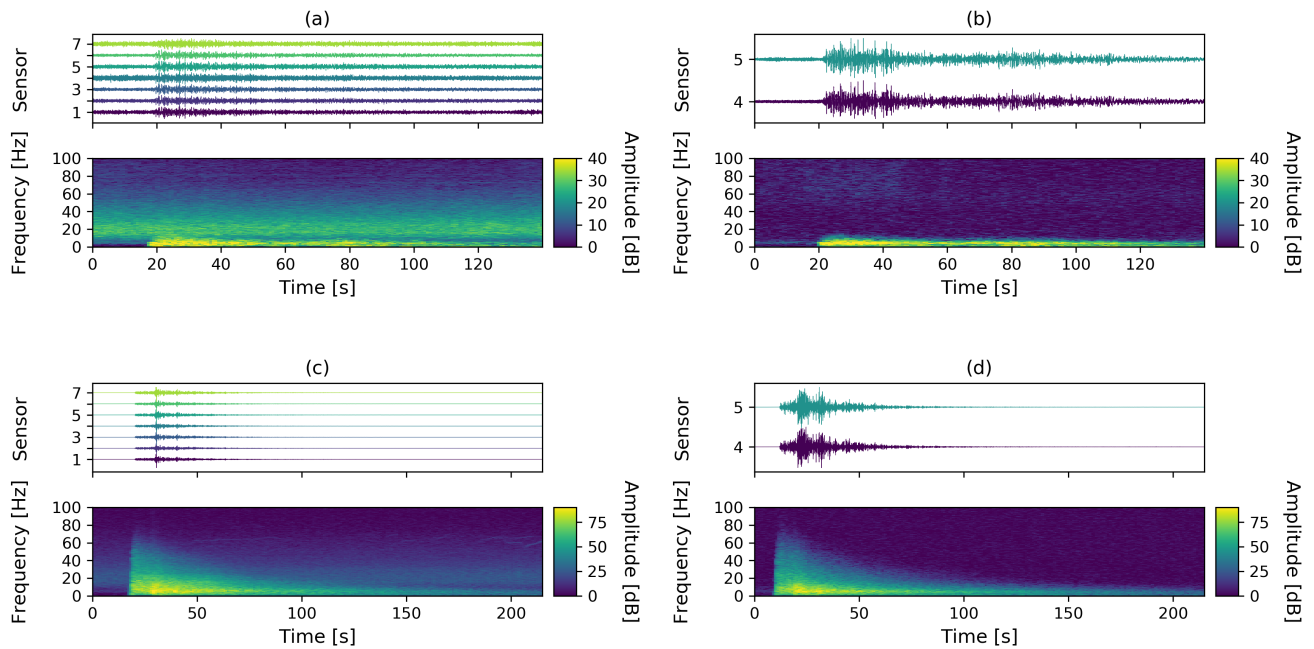

Figure 9. Time series and corresponding spectrograms for two false classifications, airplane detected at a) Dischma and b) at Wannengrat and an earthquake detected at c) Dischma and d) Wannengrat.

Figure 10. Instantaneous frequency with normalized time (normalized with event duration) for signals generated by an airplane (yellow), an earthquake (green) and an avalanche (blue). The black lines are the moving mean. The black vertical dashed lines at 0 and 1 indicate the start and the end of the events.

	single array	combined detection	location based detection
total analyzed		117	64
classified as avalanche	117	64	21
dismissed		53	43

Table 1. Number of analyzed events, events classified as an avalanche and dismissed events for the different steps in the classification workflow. For the single array classification, the total number of analyzed and dismissed events cannot be provided as the entire time series was classified.

Swiss Seismological Service (SED), we concluded that regional earthquakes within a range of $120\,km$ were recorded by both arrays. Airplanes and earthquakes were therefore recorded almost simultaneously at both arrays (Figure 9). To eliminate these false classifications, we thus used data from the Wannengrat array classified with the HMM trained with the airplane event of 31 January at 21:46 (Figure 5b). As the transient signals of airplanes and earthquakes were very similar (Figure 10), this HMM
5 also detected earthquakes.

The number of detected events at the Wannengrat array varied strongly per day (blue bars in Figure 11). 53 of these events coincided with events classified at the Dischma array and were dismissed as false detections (yellow bars in Figure 11 and Table 1). The remaining 64 potential avalanches were again mostly concentrated in January, around 9 and 10 March and early April. The distribution of these events was somewhat similar to visually observed avalanches (compare to red bars in Figure
10 7), except for the detections in April and the absence of detections at the beginning of February. Hence, we expected some misclassifications among the remaining 64 avalanche events.
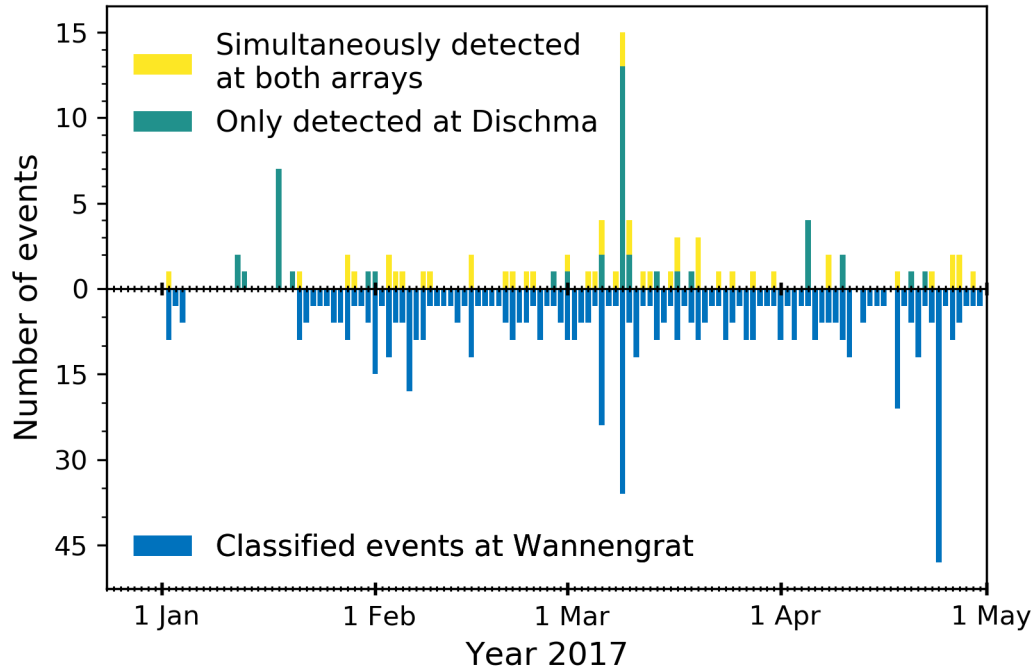
14

Figure 11. Number of events during the season. Yellow bars indicate the events detected at both arrays, the turquoise bars show the events only recorded at the Dischma array. Blue bars show the events automatically detected at the Wannengrat array. Between 5 and 20 January no data were recorded at the Wannengrat array.

4.4 Localization post-processing

We applied the MUSIC method to the remaining 64 classified events to determine the median back-azimuth paths. Events with derivatives in the median back-azimuth path larger than $10°$ were dismissed. Doing so, another 43 events were dismissed and only 21 events remained (Figure 12; Table 1). The majority of these events (i.e. 13 of 21) were observed on 9 and 10 March 2017, while the other events showed no clear link to the visual observations (red bars in Figure 12).

For each of the 21 events we determined a mean back-azimuth, which is the mean direction the signals came from. The mean back-azimuths were all pointing towards the surrounding slopes where we expected avalanches to release (Figure 13). Events with a duration longer than $100\,\mathrm{s}$ were detected coming either from the north-west or south-east.

5 Discussion

Machine learning algorithms are increasingly used for automatic signal detection in seismic data, in particular when investigating gravitational mass movements such as landslides, avalanches, rockfalls and debris flows. They include neural networks
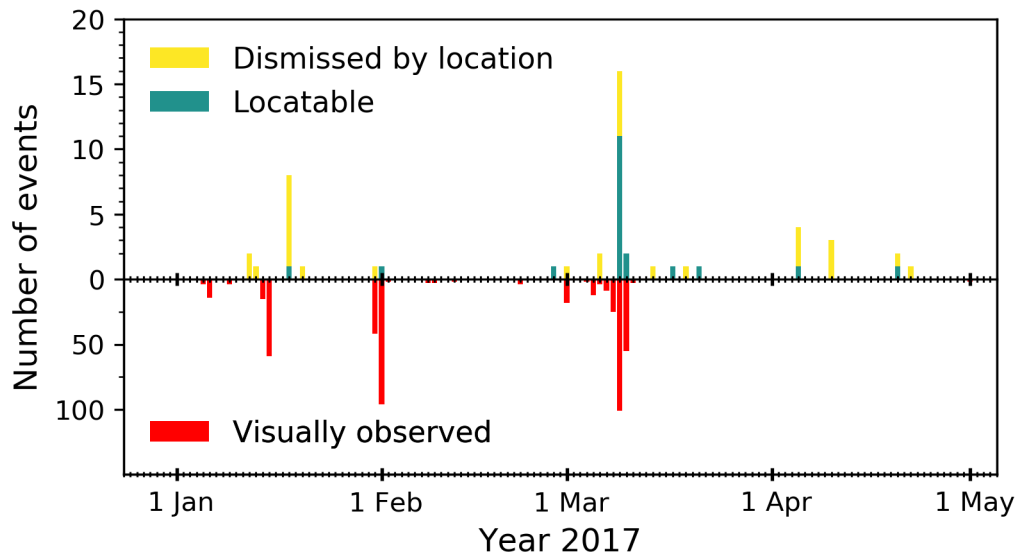
15

Figure 12. Turquoise bars are the number of events per day which are locatable and are considered as avalanches. Yellow bars are the number of events per day which could not be located and were therefore dismissed. Red bars are the number of avalanches visually observed in the area of Davos.

(Perol et al., 2018; Esposito et al., 2006), deep learning (Ross et al., 2018), random forest classifiers (Li et al., 2018; Provost et al., 2016), nearest neighbors (Bessason et al., 2007) or kurtosis-based picking (Hibert et al., 2014). While with appropriate calibration these methods generally perform rather well, the main drawback is that a large pre-labelled training database is required. The same is true for signals generated by snow avalanches (e.g. Rubin et al., 2012). The classification workflow we presented used hidden Markov models (HMMs) to automatically detect avalanches in data from seismic systems deployed above Davos, Switzerland. The approach builds on the work of Heck et al. (2018a), who adapted the HMM method developed by Hammer et al. (2017) to detect avalanches in continuous seismic data from a small aperture geophone array. A major benefit of this approach is that only one training event is required. This has substantial advantages, as the workflow could be implemented at a newly instrumented site without the need to first establish an extensive training data set. Still, a pre-operational training phase, typically one winter season, is needed to acquire at least one training event and to identify any site-specific peculiarities. Using training events recorded at different arrays might be unreliable due to possible differences in the instrumentation or heterogeneities in the local geology.

For the training event, we only used the first section of the avalanche signal, up to the maximum amplitude (Figure 5). Using the entire length of the avalanche signal resulted in poorer classification results (not shown) as there were very little variations in the transient feature behavior after the maximum in the signal. Including larger parts of the training event therefore did not provide any useful additional information for the classification.
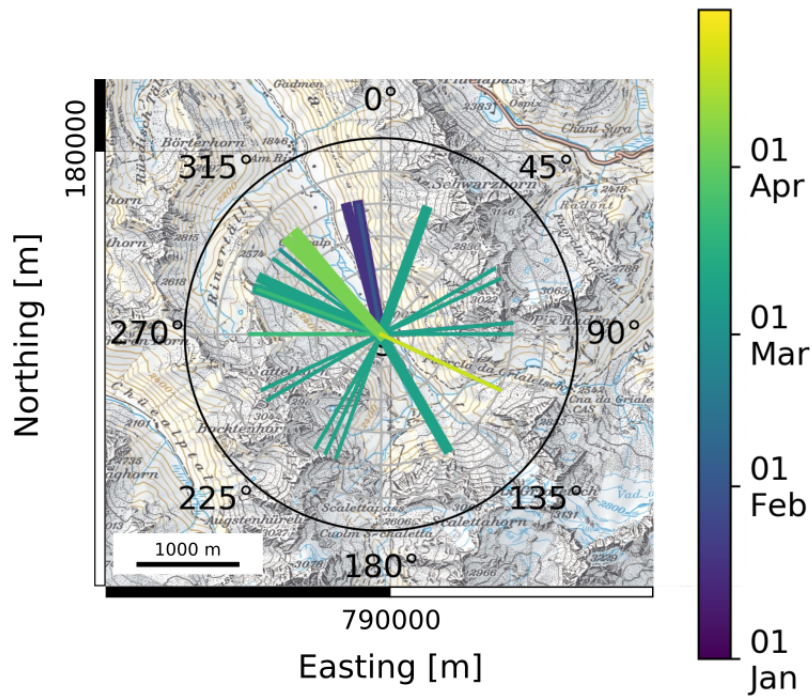
16

Figure 13. Polar plot representation overlaid on a map section of the field site. The angle represents the direction of the origin of the event, thin lines represent events with a duration $< 60\,\mathrm{s}$, thick lines events with duration $\geq 60\,\mathrm{s}$. The different colors of the lines represent the time of the year. Reproduced by permission of swisstopo (JA100118).

Heck et al. (2018a) highlighted difficulties in obtaining a reliable classifier with the HMM approach applied to a geophone array. They obtained large differences in model performance with the number of detections per sensor ranging from about 150 to over 2000. This was attributed to local heterogeneities as the sensors were packed in a styrofoam housing and inserted within the snow. In our deployment, the geophones were buried about half a meter below the ground on a flat meadow. This approach

5   resulted in a much more consistent number of detections per sensor, ranging from 125 to 169, showing that the deployment strategy can have substantial influence on the performance of the classifier.

Applying the post-processing steps suggested by Heck et al. (2018a) to remove events with a duration $\leq 12\,\mathrm{s}$ and all classifications that were classified by less than 5 sensors resulted in 117 possible avalanche events (Figure 8). Even though this approach identified the main avalanche cycle in March 2017 (compare Figure 7 and Figure 8), visual inspection of the classified

10   events indicated that at least $50\,\%$ of the events were false alarms produced by airplanes or regional earthquakes (Figure 9). Although we did not perform an exhaustive sensitivity study, some ad-hoc testing showed that these classification results did not substantially change when training a classifier with different feature combinations, changing the training event and/or the length of the training event. This highlights the difficulty in training an accurate HMM for low energy signals generated by

17

avalanches. Based on our results, we concluded that using HMMs to automatically identify avalanches in seismic data from our geophone array will inherently contain false detections, as the transient feature behavior from airplanes or regional earthquakes was very similar to signals generated by avalanches (Figure 10).

To circumvent the problem of developing an optimal event classifier for one specific array, we made use of a second array at the Wannengrat. There we performed a second classification to automatically identify airplanes and earthquakes using an event model trained with an airplane event. Since transient signals produced by earthquakes, airplanes or avalanches have similarities (Figure 10), the results obtained for the second classification based on the airplane event model likely also include avalanches and earthquakes. However, this was not a drawback since we assume that it was very unlikely that two avalanches released simultaneously at both field sites. Typically, signals generated by airplanes either have clear overtones or at least a clear Doppler effect in the signal (van Herwijnen and Schweizer, 2011a). The airplane signals that were falsely classified as avalanches with our method did not have such obvious features (Figure 9). While we do not know why and when airplanes generate such signals, and we have not identified a clear pattern explaining their presence, we have seen multiple signals like these recorded at both arrays and we are confident that these signals are generated by airplanes. Comparing the time series of detected events at both arrays allowed us to dismiss about 50% of the classified events (Figure 11). In our case, identifying co-detections across arrays is therefore an efficient approach to reduce the number of false alarms, as the weak signals generated by avalanches were only recorded at one array since the distance between both arrays was about $14\,\mathrm{km}$.

Combining the classification results from two separate arrays allowed us to reduce the number of false classifications. Nevertheless, some uncertainty remained about the origin of the identified events. In a final step, we therefore used the MUSIC method to estimate the median back-azimuth path, as suggested by Heck et al. (2018b), to further dismiss false detections. Similar approaches were suggested for the automatic detection of avalanches in infrasonic data by Marchetti et al. (2015) and Thüring et al. (2015). In those studies, the back-azimuth of continuous infrasound data was calculated using cross-correlation techniques, and only events with slight changes in back-azimuth over a predefined minimal duration were assumed as avalanche events. In contrast, we determined the back-azimuth with the MUSIC method only for those events that were already identified by the HMM, since pair-wise cross-correlation technique (beam forming) did not result in robust back-azimuth estimates for our instrumentation (Heck et al., 2018b). This last processing step further reduced the number of classified events to 21 (Figure 12).

The majority of the remaining 21 automatic classifications occurred during a period which coincided with the observed high avalanche activity in March (Figure 12). However, very few avalanches were automatically detected during the first two avalanche periods observed in the surroundings of Davos in January and February. This may be due to fact that the Dischma site is located about $12\,\mathrm{km}$ km to the southeast of Davos where avalanches are regularly observed and weather and snow conditions are sometimes different since major storms arrive from the northwest. Indeed, based on the images from the automatic cameras, very little avalanche activity was observed in the area in January and February, suggesting that there was only one main avalanche period at our site in March. With our automatic classification, it was thus possible to reconstruct the main avalanche activity period in March. Results from the localization showed that during the season avalanches released from many different slopes at the field site, in particular for the avalanche period in March (Figure 13). A seismic monitoring system

is therefore a suitable tool to monitor an area with many slopes and not just one single slope. Although the detection range is with 2 - 3 km rather limited (Heck et al., 2018b), the seismic monitoring system in combination with an automatic classifier provides great potential to identify at least the major avalanche periods. These results suggest that the detection capabilities of seismic monitoring systems are very similar to those of infrasound monitoring systems (Mayer et al., 2018).

5 Although we were able to identify one major avalanche activity period in the winter season 2016-2017, the method presented here has its limitations. Our suggested workflow requires two arrays to eliminate falsely classified events by finding co-detections. This is clearly a limiting factor as it increases the cost for the instrumentation as well as deployment and maintenance time. However, we could have directly applied the MUSIC method to all detections from the Dischma array to reduce the number of falsely classified events. Indeed, after applying the localization based step to all detections, 32 events were

10 identified as avalanches, 11 more than with the combined array and the localization based classification. Although we decided to implement a combined array classification step to save computational time, directly localizing every automatically detected event is thus also possible.

The main limitation is that we could not quantify the reliability of the classifier as no independent verification data were available. Our results suggest that the HMMs trained at the Dischma array had a false alarm rate of $\sim 80\%$, which was reduced

15 by applying the suggested post-processing steps. However, the probability of detection or the exact false alarm rate are difficult to estimate, since ground truth data are missing. As avalanches generally release during periods of poor visibility, this is a common problem when assessing the detection performance of automatic avalanche detection systems (Mayer et al., 2018; Heck et al., 2018a; Thüring et al., 2015). Alternatively, one could visually inspect the waveforms and spectrograms over the entire season (e.g. van Herwijnen et al., 2016). However, events identified in this manner often contain many uncertainties

20 (Heck et al., 2018a). For example, two of the authors independently identified possible avalanches, resulting in 44 and 23 events, respectively. Thus, the only events we could use to assess the performance of our classifier were the 13 manually identified avalanches by Heck et al. (2018b) on 9 and 10 March. Twelve of these were automatically identified with our approach, suggesting good performance in terms of probability of detection and the number of missed events. However, we cannot make any statements on the false alarm rate, which from an operational point of view is also very important.

25 Based on the approach presented here, a near real-time classification of the seismic data and hence a near real-time detection of avalanches seems possible. The computational times on a standard personal computer (8-core processor with 12 GB ram) are reasonably short as feature calculation can be performed in near real-time for all sensors simultaneously as well as the HMM construction and the classification. The MUSIC method, on the other hand, is very costly (three times real time). However, since the amount of detections for the entire season was very low, a near real-time detection could be provided with or without

30 the combined array classification. In future systems, the pre-processing steps can be integrated in the data logging unit to substantially reduce the data transmission.

6 Conclusions

During the winter season 2016-2017 we used a seismic array to continuously monitor avalanche activity in a remote area above Davos, Switzerland. By training a machine learning algorithm based on hidden Markov models (HMMs), we detected 117 events in the seismic data from January to April. Subsequent visual inspection revealed a false alarm rate of at least 50% and most of the false detections were associated with airplanes or earthquakes. We therefore trained a second HMM with data from a seismic array at a distance of $14\,km$ to remove any co-detections. Finally, we applied a multiple signal classification (MUSIC) approach to define threshold criteria for automatic avalanche identification by considering avalanches as a moving source. Overall, this workflow resulted in 21 automatic classifications. The majority of these events occurred on 9 and 10 March 2017, in accordance with a period of high avalanche activity observed in the surroundings of Davos.

Our results show that seismic monitoring systems in combination with an automatic classifier provides great potential to identify at least the major avalanche periods. In our workflow, using an array processing method to determine the source of the seismic events was of crucial importance to reduce falsely classified events. In future experiments we want to introduce an additional array within a shorter distance to improve the localization and no longer require the combined array classification approach.

Competing interests. The authors declare that they have no conflict of interest.

References

Barnes, A. E.: Instantaneous spectral bandwidth and dominant frequency with applications to seismic reflection data, Geophysics, 58, 419–428, 1993.

Bessason, B., Eiriksson, G., Thorarinsson, O., Thorarinsson, A., and Einarsson, S.: Automatic detection of avalanches and debris flows by seismic methods, J. Glaciol., 53, 461–472, https://doi.org/10.3189/002214307783258468, 2007.

Beyreuther, M., Hammer, C., Wassermann, J., Ohrnberger, M., and Megies, T.: Constructing a Hidden Markov Model based earthquake detector: application to induced seismicity, Geophys. J. Int., 189, 602–610, https://doi.org/10.1111/j.1365-246X.2012.05361.x, 2012.

Esposito, A. M., Giudicepietro, F., Scarpetta, S., D'Auria, L., Marinaro, M., and Martini, M.: Automatic discrimination among landslide, explosion-quake and microtremor seismic signals at Stromboli Volcano using neural networks, Bull. Seismol. Soc. Am., 96, 1230, https://doi.org/10.1785/0120050097, +http://dx.doi.org/10.1785/0120050097, 2006.

Hammer, C., Beyreuther, M., and Ohrnberger, M.: A seismic-event spotting system for Volcano fast-response systems, Bull. Seismol. Soc. Am., 102, 948–960, https://doi.org/10.1785/0120110167, 2012.

Hammer, C., Ohrnberger, M., and Fäh, D.: Classifying seismic waveforms from scratch: a case study in the alpine environment, Geophys. J. Int., 192, 425–439, 2013.

Hammer, C., Fäh, D., and Ohrnberger, M.: Automatic detection of wet-snow avalanche seismic signals, Nat. Hazards, 86, 601–618, https://doi.org/10.1007/s11069-016-2707-0, 2017.

Heck, M., Hammer, C., van Herwijnen, A., Schweizer, J., and Fäh, D.: Automatic detection of snow avalanches in continuous seismic data using hidden Markov models, Nat. Hazards Earth Syst. Sci., 18, 383–396, https://doi.org/10.5194/nhess-18-383-2018, 2018a.

Heck, M., Hobiger, M., van Herwijnen, A., Schweizer, J., and Fäh, D.: Localization of seismic events produced by avalanches using multiple signal classifications, Geophys. J. Int., https://doi.org/10.1093/gji/ggy394, http://dx.doi.org/10.1093/gji/ggy394, 2018b.

Hibert, C., Mangeney, A., Grandjean, G., Baillard, C., Rivet, D., Shapiro, N. M., Satriano, C., Maggi, A., Boissier, P., Ferrazzini, V., and Crawford, W.: Automated identification, location, and volume estimation of rockfalls at Piton de la Fournaise volcano, Journal of Geophysical Research: Earth Surface, 119, 1082–1105, https://doi.org/10.1002/2013JF002970, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2013JF002970, 2014.

Joswig, M.: Knowledge-Based Seismogram Processing by Mental Images, IEEE Trans. Syst. Man. Cybern., 24, 429 – 439, 1994.

Kanasewich, E. R.: Time Sequence Analysis in Geophysics, The University of Alberta Press, Edmonton, Alberta, Canada, 1981.

Kramer, S.: Geotechnical Earthquake Engineering, Always learning, Prentice-Hall international series in Civil Engineering and engineering mechanics, Upper Saddle River, New Jersey, 1996.

Lacroix, P., Grasso, J.-R., Roulle, J., Giraud, G., Goetz, D., Morin, S., and Helmstetter, A.: Monitoring of snow avalanches using a seismic array: Location, speed estimation, and relationships to meteorological variables, J. Geophys. Res. Earth Surf, 117, F01 034, https://doi.org/10.1029/2011JF002106, 2012.

Leprettre, B., Navarre, J., and Taillefer, A.: First results from a pre-operational system for automatic detection and recognition of seismic signals associated with avalanches, J. Glaciol., 42, 352–363, https://doi.org/10.1109/ICMLA.2012.12, 1996.

Li, Z., Meier, M.-A., Hauksson, E., Zhan, Z., and Andrews, J.: Machine Learning Seismic Wave Discrimination: Application to Earthquake Early Warning, Geophysical Research Letters, 45, 4773–4779, https://doi.org/10.1029/2018GL077870, 2018.

Marchetti, E., Ripepe, M., Ulivieri, G., and Kogelnig, A.: Infrasound array criteria for automatic detection and front velocity estimation of snow avalanches: towards a real-time early-warning system, Nat. Hazards Earth Syst. Sci., 15, 2545 – 2555, https://doi.org/10.5194/nhess-15-2545-2015, 2015.

5   Mayer, S., van Herwijnen, A., Ulivieri, G., and Schweizer, J.: Evaluating the performance of an operational infrasound avalanche detection system at three locations in the Swiss Alps during two winter seasons, in: Proc. 2018 Int. Snow Sci. Workshop, Innsbruck, Austria, 7 - 12 October 2018, 2018.

McClung, D. and Schaerer, P. A.: The Avalanche Handbook, The Mountaineers Books, 2006.

Nishimura, K. and Izumi, K.: Seismic signals induced by snow avalanche flow, Nat. Hazards, 15, 89–100, https://doi.org/10.1023/A:1007934815584, 1997.

10   Ohrnberger, M.: Continuous automatic classification of seismic signals of volcanic origin at Mt. Merapi, Java, Indonesia, PhD. thesis, 2001.

Perol, T., Gharbi, M., and Denolle, M.: Convolutional neural network for earthquake detection and location, Science Advances, 4, https://doi.org/10.1126/sciadv.1700578, 2018.

Provost, F., Hibert, C., and Malet, J.-P.: Automatic classification of endogenous landslide seismicity using the Random Forest supervised classifier, Geophysical Research Letters, 44, 113–120, https://doi.org/10.1002/2016GL070709, 2016.

15   Ross, Z. E., Meier, M.-A., and Hauksson, E.: P Wave Arrival Picking and First-Motion Polarity Determination With Deep Learning, Journal of Geophysical Research: Solid Earth, 123, 5120–5129, https://doi.org/10.1029/2017JB015251, 2018.

Rost, S. and Thomas, C.: Array seismology: methods and application, Rev. Geophys., 40, 2.1–2.27, https://doi.org/10.1029/2000RG000100, 2002.

Rubin, M., Camp, T., van Herwijnen, A., and Schweizer, J.: Automatically detecting avalanche events in passive seismic data, IEEE International Conference on Machine Learning and Applications, 1, 13–20, https://doi.org/10.1109/ICMLA.2012.12, 2012.

20

Schaerer, P. A. and Salway, A. A.: Seismic and impact-pressure monitoring of flowing avalanches, J. Glaciol., 26, 179–187, https://doi.org/10.1017/S0022143000010716, 1980.

Schmidt, R.: Multiple emitter location and signal parameter estimation, IEEE Trans. Antennas Propag., 34, 276–280, https://doi.org/10.1109/TAP.1986.1143830, 1986.

25   Scott, E., Hayward, C., Kubichek, R., Hamann, J., Pierre, J., Comey, B., and Mendenhall, T.: Single and multiple sensor identification of avalanche-generated infrasound, Cold Reg. Sci. Technol., 47, 159–170, https://doi.org/10.1016/j.coldregions.2006.08.005, 2007.

Steinkogler, W., Meier, L., Langeland, S., and Wyssen, S.: Operational radar and infrasound systems for avalanche detection, in: Proc. 2016 Int. Snow Sci. Workshop, Breckenridge, Colorado, USA, 3 - 7 October 2016, pp. 309–315, 2016.

Suriñach, E., Furdada, G., Sabot, F., Biescas, B., and Vilaplana, J.: On the characterization of seismic signals generated by snow avalanches for monitoring purposes, Ann. Glaciol., 32, 268–274, https://doi.org/10.3189/172756401781819634, 2001.

30

Suriñach, E., Vilajosana, I., Khazaradze, G., Biescas, B., Furdada, G., and Vilaplana, J.: Seismic detection and characterization of landslides and other mass movements, Nat. Hazards Earth Syst. Sci., 5, 791–798, https://doi.org/10.5194/nhess-5-791-2005, 2005.

Taner, M. T., Koehler, F., and Sheriff, R.: Complex seismic trace analysis, Geophysics, 44, 1041–1063, 1979.

Thüring, T., Schoch, M., van Herwijnen, A., and Schweizer, J.: Robust snow avalanche detection using supervised machine learning with infrasonic sensor arrays, Cold Reg. Sci. Tech., 111, 60 – 66, https://doi.org/10.1016/j.coldregions.2014.12.014, 2015.

35

van Herwijnen, A. and Schweizer, J.: Seismic sensor array for monitoring an avalanche start zone: design, deployment and preliminary results, J. Glaciol., 57, 257–264, https://doi.org/10.3189/002214311796405933, 2011a.

van Herwijnen, A. and Schweizer, J.: Monitoring avalanche activity using a seismic sensor, Cold Reg. Sci. Tech., 69, 165–176, https://doi.org/10.1016/j.coldregions.2011.06.008, 2011b.

van Herwijnen, A., Heck, M., and Schweizer, J.: Forecasting snow avalanches by using highly resolved avalanche activity data obtained through seismic monitoring, Cold Reg. Sci. Tech., 132, 68–80, https://doi.org/10.1016/j.coldregions.2016.09.014, 2016.