TOWARDS THE AUTOMATIC DETECTION OF AVALANCHES IN SEISMIC DATA USING HIDDEN MARKOV MODELS

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

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

ABSTRACT: Seismic monitoring systems are well suited for the remote detection of various mass movements, including snow avalanches. While this has been known since the 1970s, thus far seismic monitoring systems are not widely used in operational forecasting. The difficulty lies in automatically identifying seismic signals generated by avalanches in continuous seismic data streams. While a few studies suggested methods to resolve this problem, high false alarm rates or a lack of validation generally hindered their implementation. We therefore aim at automatically detecting avalanches in the continuous seismic data stream and evaluate model performance in terms of probability of detection and false alarm rate. By using a Hidden Markov Model (HMM), a statistical pattern recognition tool widely used for speech recognition, we construct a wide background. Then we learn from only one single avalanche sample, an event specific classifier to automatically detect avalanches. Using seismic data obtained from avalanche start zones above Davos from 2010, we compare classified events with avalanches manually identified in the seismic data. Using one single classifier, as was done in previous studies using HMM, results in unsatisfactory model performance. This is largely due to continuous changes in environmental noise. Regularly updating the background model is therefore required, resulting in improved overall model performance. While substantial progress is still required, overall our results suggest that the automatic detection of avalanches in seismic data is feasible, bringing us one step closer to implementing seismic monitoring systems in operational forecasting.

KEYWORDS: automatic detection, classification, seismic monitoring

1. INTRODUCTION

Data on avalanche activity is of crucial importance for avalanche forecasting services. However, such data are hard to come by, especially for remote areas and during periods of poor visibility (e.g. snowfall, fog and at night). To resolve this issue, we investigate a method to automatically detect avalanches in continuous data obtained from seismic monitoring systems.

Seismic systems are well suited to detect mass movements such as rock falls, pyroclastic flows as well as snow and ice avalanches (Caplan-Auerbach et al. (2004); Suriñach et al. (2005)). First experiments on monitoring avalanches using seismic sensors were conducted in the 1970s (St. Lawrence and Williams (1976) and Harrison (1976)), identifying the characteristic spindle shape of the seismogram generated by avalanches. Subsequent research has shown that the frequency spectrum and the spectrogram (i.e. a time-frequency representation of seismic data) of seismic signals generated by avalanches are different from other known seismic sources (Kishimura and Izumi (1997) and Sabot et al. (1998)). Although seismic signals generated by avalanches have distinct characteristics, only little effort has been made to automatically detect avalanches using machine-learning algorithms.

Leprettre et al. (1996) were the first to attempt this feat using fuzzy logic sets, while more recently Bessason et al. (2007) used a K-means nearest neighbor method. Although both methods were able to automatically detect avalanches, the number of false alarms was high. Therefore, Rubin et al. (2012) compared the performance of twelve modern machine-learning algorithms. While all methods were able to detect most avalanches, the number of false alarms was too high for operational purposes.

Hidden Markov models HMMs are well suited to automatically classify seismic data from monitoring systems and have been used to, for instance, construct a volcano fast response system (Hammer et al. (2012)), an earthquake detector (Beyreuther et al. (2012)) or to detect rockfalls, earthquakes and quarry blasts on seismic broad-

* *Corresponding author address:*
Matthias Heck, WSL Institute for Snow and Avalanche Research SLF, Flüelastrasse 11, CH-7260 Davos, Switzerland
email: heck@slf.ch

band stations of the Swiss seismological service (Hammer et al. (2013), Dammeier et al. (2015)). Simple algorithms, which only take the distribution of features of the seismic signals into account, e.g. cluster algorithms, would probably classify different events as the same, if they have a similar distribution of the features. However, the advantages of HMMs compared to those simple algorithms are, that the time dependent behavior of the signals is also taken into account. A hidden Markov model has the power to identify events with a similar behavior and classify them in the same class. In order to decrease the dependency on previous acquired training samples we apply a novel training approach to construct corresponding HMMs (Hammer et al. (2012)). In this approach the whole data set is used to create a background model, so information about all possible events is already within the classifier. In the following, a small amount of training samples is sufficient to learn the classifier.
Our aim is, to use a HMM to automatically detect avalanches in a continuous seismic waveform.

2. METHODS

2.1 Instumentation and field site

The seismic data we analyzed were obtained using a geophone array deployed at the Steintälli field site above Davos, Switzerland. The seismic array was deployed on a steep NE facing slope at 2500 m.a.s.l. and surrounded by various avalanche starting zones (Figure 1).
The seismic array consisted of seven vertical geophones with an eigenfrequency of 14 Hz. The sensors were inserted in the snow.
A low power 24-bit data acquisition system was used to continuously acquire data from the sensors at a sampling rate of 500 Hz. Data were stored locally on a low power computer and manually retrieved approximately every 10 days.
More details about the field site, sensor array and data acquisition system can be found in van Herwijnen et al. (2014); van Herwijnen and Schweizer (2011a); van Herwijnen and Schweizer (2011b); and van Herwijnen et al. (2010).

Fig 1: Picture of the Steintälli field site. The geophones are deployed within the red circle.

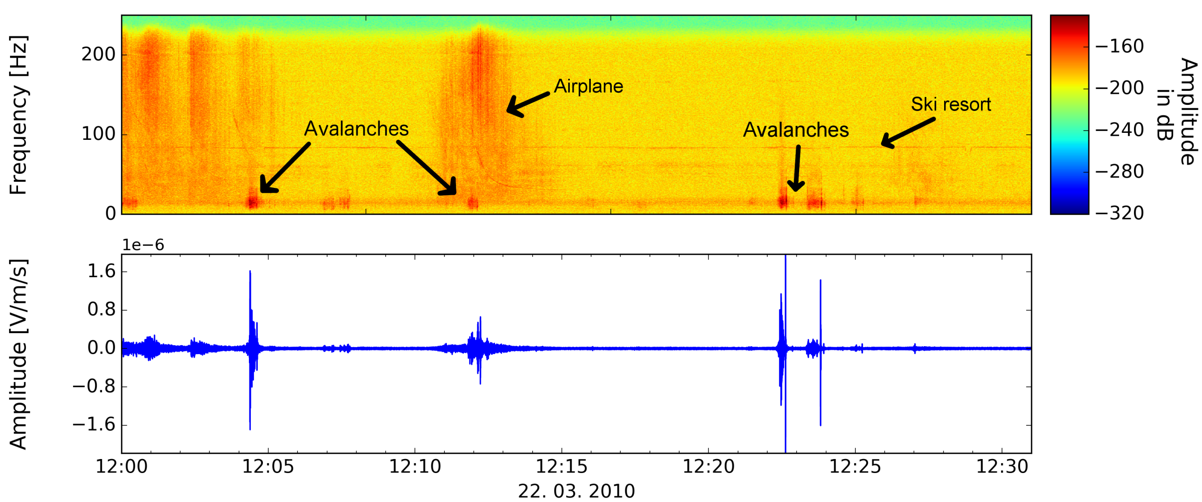

Fig 2: Bottom: Part of the continuous waveform containing avalanches and airplane events. Top: corresponding seismogram

2.2 *Data*

We used continuous seismic data acquired with the seismic array from mid January to end of April 2010. Due to the low frequency of the geophones and the high sensitivity of the data acquisition system, there was considerable background noise in the data. Indeed, as can be seen in Figure 2, apart from avalanches, there are various sources, which generate distinct seismic signals. It is not possible to distinguish these signals by simple visual observations of the time series (bottom in Figure 2). However, when looking at the spectrogram (top in Figure 2), the various signals become apparent. Based on images obtained from automatic cameras, field observations and their personal experience, van Herwijnen and Schweizer (2011a) performed a visual analysis of the seismic data and were able to identify 385 avalanches between 12 January and 30 April 2010 (Figure 3). Of these avalanches, 33 were confirmed through visual observations, i.e. from field observations or by inspecting images taken with an automatic camera. Overall, there were two distinct peaks in the avalanche activity; the first around 22 March and the second around 24 April 2010.
We used this avalanche catalogue as our reference data set for the automatic classification of the continuous seismic data.

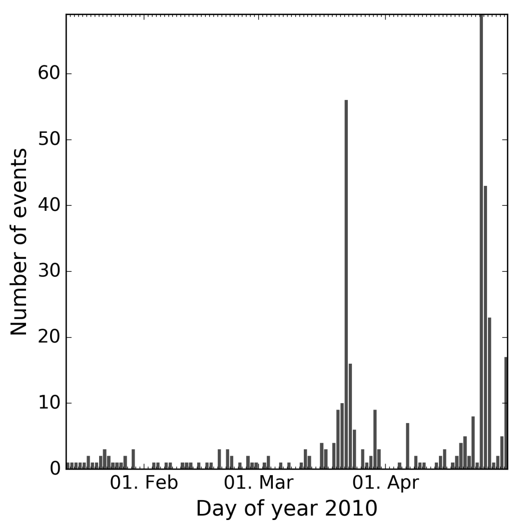

Fig 3: Number of avalanches per day between 12 January and 30 April 2010

2.3 *Hidden Markov models*

Hidden Markov models (HMMs) are statistical pattern recognition models commonly used, for instance, for speech recognition tasks (e.g. Rabiner (1989)). In order to decrease the redundant information a compressed representation of the data, acts as input for the classifier. Short-term wavefield attributes, so-called features, replace the raw seismic waveform.

Features are calculated by applying a sliding window to the seismic data. For example, in Figure 4 the central frequency of a seismic event is calculated for each time window (red bars). By defining the increment and the overlap of the windows, the density of calculated values can be defined. This example clearly shows that the central frequency of noise is around 100 Hz. When the event (avalanche) begins, the central frequency decreases to 20 Hz until the event is over and the central frequency returns to a value around 100 Hz. By combining various features, different types of events are represented in a unique way. In total we calculated 19 features, including the central frequency, the bandwidth and the dominant frequency. A list of all the features and the corresponding functions can be found in Hammer et al. (2012).

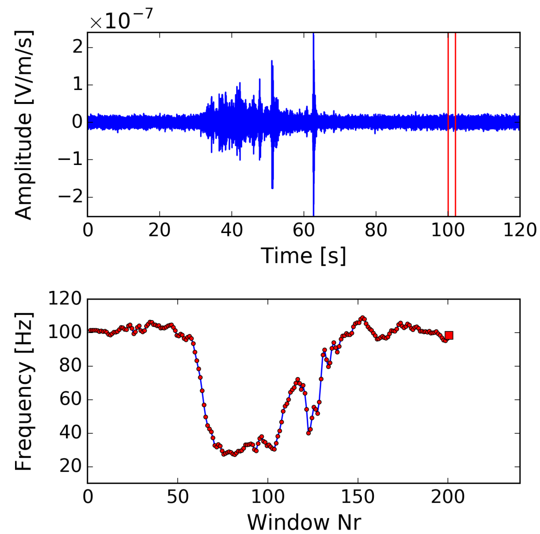

Fig 4: Example for the feature calculation. On top the seismic trace. The red bars indicate the width of the window used to calculate a feature value. Bottom: central frequency with window number of the seismic trace.

To create a hidden Markov model, the statistical distribution of the features is important. In the used approach, a 19-dimensional probability distribution for each feature consisting of 16 Gaussians is constructed, which builds the so-called background model. To create this background model it is important to include many different types of events since the sensitivity of the model increases if more information is contained in the background model. Once the background model is created, the event-specific HMM has to be trained. Therefore we used one event listed in our reference data set. Since there are different types of avalanches (dry- and wet-snow) listed, we compared the features of several events. Results showed that the seismic signatures remained the same, independent of the type of avalanche. One single avalanche class is thus sufficient for our classification task. We therefore used an avalanche recorded on 22 March 2010 at 13:44 (Figure 5) as the training event.

3. RESULTS

The amount of data used to create the background model is of vital importance for the performance. Too few data would make the model 'blind' for rare events and too much data would increase the computational time resulting in an impractical method. We therefore investigated the influence of the length of the background model on the classification performance.

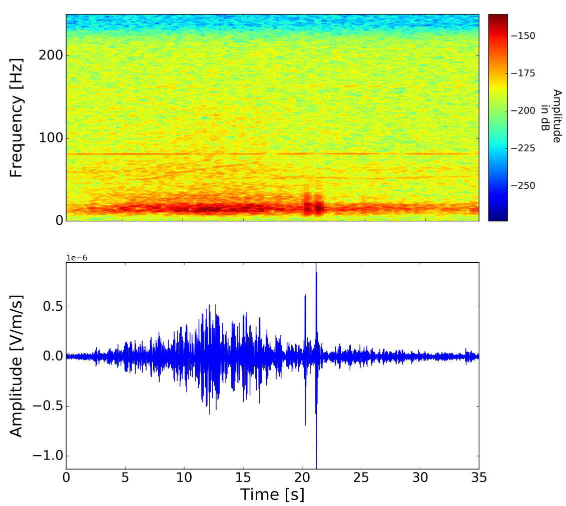

Fig 5: Bottom: Seismic trace of the avalanche training event used for the hidden Markov model recorded on the 22nd March 2010 Top: corresponding spectrogram

3.1 *Single day static background model*

For this first iteration we used data recorded on one day in January to create the background model and then classified the entire winter.

The two periods of high avalanche activity were relatively well captured as the algorithm detected a total amount of 120 events around the 22 March and 102 events around the 24 April. However, the overall avalanche activity pattern did not resemble the reference data (compare Figures 3 and 6) due to many false alarms, e.g. around 10 February. While the performance of this single day background model was reasonable in terms of probability of detection (POD = 44.2%), the very high false alarm rate (FAR = 88.8%) makes such a model unusable.

3.2 *Adaptive background model*

The results using a single model for the classification were not satisfying. This is because for our seismic array the feature distribution was not constant in time. Indeed, diurnal and seasonal variations were present in the features (not shown). A single model is therefore not suitable to classify an entire winter season, and we investigated the influence of recalculating the background model for a predefined time interval Δt, ranging from one to 24 hours. With this method the goal was to account for the variability of seismic signals to improve the overall classification.

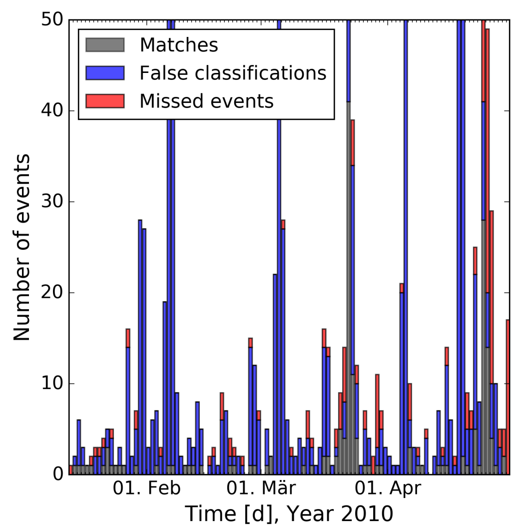

Fig 6: Avalanche activity data obtained using a one-day background model to classify the entire season showing the number of correctly classified avalanches (gray bars), the false classifications (blue bars) and the missed events (red bars).

For each time interval Δt, we recalculated a new background model and used it to classify the data within Δt. We used the entire data within Δt to construct the background model to provide as much information as possible to the classifier and relied on all 19 features. The training event was the same sample event as we used for the first classification described above.

Overall, when using an adaptive background model, far fewer events were classified as avalanche than with a static background model. This resulted in lower false alarm rates, but also fewer avalanches were correctly classified resulting in lower POD values (Table 1). However, POD values increased substantially with increasing time interval, from 8 to 28%, while the FAR values remained mostly the same.
The best performance was obtained for longer time intervals of 12 and 24 hours (Table 1). With a 24-hour model, the two avalanche cycles in March and April were well captured, and the overall avalanche activity data more closely resembled the reference data (compare Figures 3 and 7).

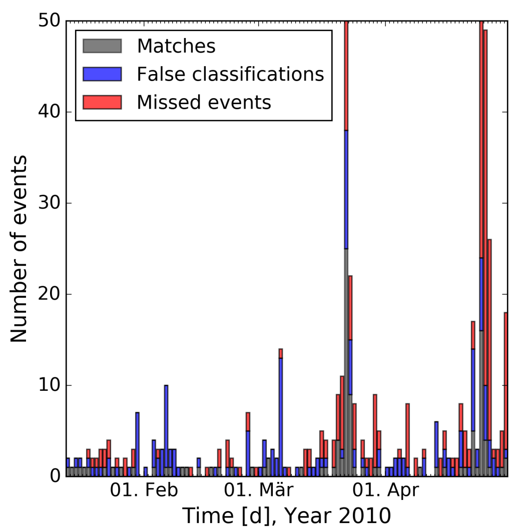

Fig7: Avalanche activity data obtained using a 24-hour adaptive background model showing the number of correctly classified avalanches (gray bars), the false classifications (blue bars) and the missed events (red bars).

Tbl 1: Probability of detection (POD) and false alarm rate (FAR) for the automatic classification of avalanches using a HMM with and adaptive background model updated every time interval Δt to classify the data within the same time interval.

Time interval Δt	*No of detected events*	*POD*	*FAR*
1h	32	8.3%	56.8%
2h	47	12.2%	53.9%
6h	80	20.7%	52.7%
12h	92	23.8%	56.4%
24h	106	27.5%	60.9%

3.3 Operational application

In the previous classifications, we built a background model using a predefined amount of data to classify the same data. This can be done to verify of HMM can be used to automatically detect avalanches. However, such models are impractical for operational applications because there is always a lag between the time the data is gathered and the time it is analyzed.

To overcome this flaw, we further investigated an operational workflow, in which the data already recorded within a time interval Δt are used to create the background model and incoming data within the time interval Δt is classified on the fly (Figure 8).

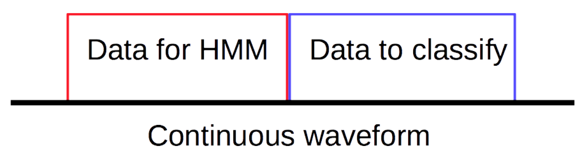

Fig 8: Schematic representation of the operational workflow. Past data within the red time interval is used to create the background model to classify data within the blue time interval on the fly.

As before, we compared several time intervals in order to investigate the best suited length for operational purposes.

Tbl 2: Probability of detection (POD) and false alarm rate (FAR) for the automatic classification of avalanches using an operational HMM with a background model constructed with past data within a time interval of Δt to classify incoming data within a time interval of Δt.

Time interval Δt	*Nr of detected events*	*POD*	*FAR*
1h	140	35.5%	88.8%
3h	134	34.6%	90.3%
6h	149	38.4%	92.2%
12h	149	38.6%	85.8%
24h	162	42%	87.4%

Again, the longer time intervals performed best. However, as with the static model, the false alarm rates were very high (≥ 85%), making such models impractical.

4. CONCLUSIONS AND OUTLOOK

We implemented a sophisticated machine-learning algorithm, a hidden Markov model (HMM), for the automatic detection of avalanches in continuous seismic data. Although HMM have been shown to perform well in previous classification tasks (e.g. Hammer et al., 2012), the automatic detection of avalanches in seismic data obtained by our geophone array remains relatively poor.
Indeed, our results show that one single model created at the beginning of the season is not capable of automatically detecting avalanches with a low false alarm rate. Although using a more adaptive model, by periodically updating the background model resulted in lower false alarm rates; both the probability of detection and the false alarm rate remained relatively poor.

There are several reasons for the poor performance. First, the reference avalanche catalogue used to evaluate model performance was obtained by visual inspection of the seismic traces and their corresponding spectrograms. This database is therefore somewhat biased. Furthermore, it mostly contains signals generated by rather small avalanches (van Herwijnen et al., 2016). With this in mind, the actual performance of the model is likely underestimated.

Second, we used completely different seismic instrumentation than in other studies. Indeed, in previous studies low frequency broadband seismometers were used (e.g. Hammer et al. (2013)). These sensors are coupled to the ground and often installed in underground caverns where there is little environmental noise. Our instrumentation consisted of low cost high frequency geophones inserted in the snow. Thus, there was only a thin snow cover shielding our sensors. This resulted in substantial environmental noise in our seismic data, which varied at time scales ranging from hours to months. These variable background noise levels significantly complicated the classification task, as highlighted by the relatively poor model performance.

Although the results presented here might suggest that a HMM is not very well suited for the automatic detection of avalanches in continuous seismic data, we believe that there is a lot of room for improvement. From an instrumentation point of view, we have replaced our sensors with lower frequency (4.5 Hz rather than 14 Hz) geophones and now attach them directly to rocks in the ground. This will increase the signal-to-noise ratio, which should result in a better performance. From a signal processing point of view, future efforts to improve the model will have to take into account signal duration, which is an indirect measure of avalanche size (van Herwijnen et al., 2013). Furthermore, there are many more parameters that can be optimized to improve model performance, including downsampling the data, using fewer features and using different time intervals for building the background model and for classifying the data.

Overall, our results confirm that automatically classifying avalanches in seismic data is not an easy task. Nevertheless, while substantial progress is still required, overall our results suggest that the automatic detection of avalanches in seismic data using HMM is feasible, bringing us one (small) step closer to implementing seismic monitoring systems in operational forecasting

REFERENCES

Bessason, B., Eiriksson, G., Thorarinsson, O., Thorarinsson, A. and Einarsson, S., 2007. Automatic detection of avalanches and debris flows by seismic methods. J. Glaciol., 53(182): 461-472.

Beyreuther, M., Hammer, C., Wassermann, J., Ohrnberger, M. and Megies, T., 2012. Constructing a Hidden Markov Model based earthquake detector: application to induced seismicity. Geophys. J. Int., 189(1): 602-610.

Caplan-Auerbach, J. and Huggel, C., 2007. Precursory seismicity associated with frequent, large ice avalanches on Iliamna volcano, Alaska, USA. J. Glaciol., 53(180): 128-140.

Dammeier, F, J.F. Moore JR, C. Hammer, F. Haslinger and S. Loew, 2016: Automatic detection of alpine rockslides in continuous seismic data using hidden Markov models. J. Geophys. Res. Earth 121 (2), 351–371

Hammer, C., Beyreuther, M. and Ohrnberger, M., 2012. A seismic-event spotting system for volcano fast-response systems. B. Seismol. Soc. Am., 102(3): 948-960.

Hammer, C., Ohrnberger, M. and Fäh, D., 2013. Classifying seismic waveforms from scratch: a case study in the alpine environment. Geophys. J. Int., 192(1): 425-439.

Harrison, J.C., 1976. Seismic signals from avalanches. In: R. Armstrong and J.D. Ives (Editors), Avalanche release and snow characteristics, San Juan Mountains, Colorado. Sun Juan Avalanche Project, Final Report 1971-1975, Occasional Paper No. 19. Institute of Arctic and Alpine Research, University of Colorado, Boulder CO, U.S.A., pp. 145-150.

Leprettre, B.J.P., Navarre, J.P. and Taillefer, A., 1996. First results from a pre-operational system for automatic detection and recognition of seismic signals associated with avalanches. J. Glaciol., 42(141): 352-363.

Nishimura, K. and Izumi, K., 1997. Seismic signals induced by snow avalanche flow. Nat. Hazards, 15(1): 89-100.

Rabiner, L.R., 1989. A tutorial on Hidden Markov Models and selected application in speech recognition. Proceedings of the IEEE, 77(2): 257-286.

Rubin, M.J., Camp, T., van Herwijnen, A. and Schweizer, J., 2012. Automatically detecting avalanche events in passive seismic data. In: M. Arif Wani, T. Khoshgoftaar, X. Zhu and N. Seliya (Editors), 2012 Eleventh International Conference on Machine Learning and Applications (ICMLA), pp. 13-20.

Sabot, F., Naaim, M., Granada, F., Suriñach, E., Planet, P. and Furdada, G., 1998. Study of avalanche dynamics by seismic methods, image-processing techniques and numerical models. Ann. Glaciol., 26: 319–323.

Schweizer, J. and van Herwijnen, A., 2013. Can near real-time avalanche occurrence data improve avalanche forecasting? In: F. Naaim-Bouvet, Y. Durand and R. Lambert (Editors), Proceedings ISSW 2013. International Snow Science Workshop, Grenoble, France, 7-11 October 2013. ANENA, IRSTEA, Météo-France, Grenoble, France, pp. 195-198.

St.Lawrence, W.F. and Williams, T.R., 1976. Seismic signals associated with avalanches. J. Glaciol., 17(77): 521-526.

Suriñach, E., Vilajosana, I., Khazaradze, G., Biescas, B., Furdada, G. and Vilaplana, J.M., 2005. Seismic detection and characterization of landslides and other mass movements. Nat. Hazards Earth Syst. Sci., 5(6): 791-798.

van Herwijnen, A., Dreier, L. and Bartelt, P., 2013. Towards a basic avalanche characterization based on the generated seismic signal. In: F. Naaim-Bouvet, Y. Durand and R. Lambert (Editors), Proceedings ISSW 2013. International Snow Science Workshop, Grenoble, France, 7-11 October 2013. ANENA, IRSTEA, Météo-France, Grenoble, France, pp. 1033-1037.

van Herwijnen, A., Heck, M. and Schweizer, J., 2016. Forecasting snow avalanches by using avalanche activity data obtained through seismic monitoring. Cold Reg. Sci. Technol.: submitted.

van Herwijnen, A. and Schweizer, J., 2011a. Monitoring avalanche activity using a seismic sensor. Cold Reg. Sci. Technol., 69(2-3): 165-176.

van Herwijnen, A. and Schweizer, J., 2011b. Seismic sensor array for monitoring an avalanche start zone: design, deployment and preliminary results. J. Glaciol., 57(202): 267-276.