



## UvA-DARE (Digital Academic Repository)

### Crack detection in earth dam and levee passive seismic data using support vector machines

Fisher, W.D.; Camp, T.K.; Krzhizhanovskaya, V.V.

**DOI**

[10.1016/j.procs.2016.05.339](https://doi.org/10.1016/j.procs.2016.05.339)

**Publication date**

2016

**Document Version**

Final published version

**Published in**

Procedia Computer Science

**License**

CC BY-NC-ND

[Link to publication](#)

**Citation for published version (APA):**

Fisher, W. D., Camp, T. K., & Krzhizhanovskaya, V. V. (2016). Crack detection in earth dam and levee passive seismic data using support vector machines. *Procedia Computer Science*, 80, 577-586. <https://doi.org/10.1016/j.procs.2016.05.339>

**General rights**

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

**Disclaimer/Complaints regulations**

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

*UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)*



# Crack Detection in Earth Dam and Levee Passive Seismic Data Using Support Vector Machines

Wendy D. Fisher<sup>1</sup>, Tracy K. Camp<sup>1</sup>, and Valeria V. Krzhizhanovskaya<sup>2,3,4</sup>

<sup>1</sup> Colorado School of Mines, Golden, Colorado, U.S.A.  
wbelcher@mines.edu, tcamp@mines.edu

<sup>2</sup> University of Amsterdam, Science Park 904, 1098 XH, Amsterdam, the Netherlands  
V.Krzhizhanovskaya@uva.nl

<sup>3</sup> National Research University ITMO, St. Petersburg, 197101, Russia

<sup>4</sup> St. Petersburg State Polytechnic University, St. Petersburg, 195251, Russia

## Abstract

We investigate techniques for earth dam and levee health monitoring and automatic detection of anomalous events in passive seismic data. We have developed a novel data-driven workflow that uses machine learning and geophysical data collected from sensors located on the surface of the levee to identify internal erosion events. In this paper, we describe our research experiments with binary and one-class Support Vector Machines (SVMs). We used experimental data from a laboratory earth embankment (80% normal and 20% anomalies) and extracted nine spectral features from decomposed segments of the time series data. The two-class SVM with 10-fold cross validation achieved over 97% accuracy. Experiments with the one-class SVM use the top two features selected by the ReliefF algorithm and our results show that we can successfully separate normal from anomalous data observations with over 83% accuracy.

*Keywords:* Data-driven levee monitoring, machine learning, anomaly detection, passive seismic.

## 1 Introduction

In this paper, we describe our research for the advancement of earth levee health assessment. We are developing a novel data-driven workflow for the automatic detection of anomalous events that uses machine learning and geophysical data to identify internal erosion events. Our lightweight anomaly detection scheme builds upon our work using unsupervised clustering [1], which shows a clear separation of events (e.g., cracks) from non-events. We begin by discussing the background and motivation for our application to earth levee passive seismic data.

### 1.1 Identifying Internal Erosion Events in Earth Dams and Levees

Earth dams and levees are constructed with earthen materials such as rock, sand, and clay [2] and are built primarily for flood control, water storage, and irrigation. The main causes of earth

levee failures are typically due to piping, slope instability, foundation issues, or overtopping [3]. Figure 1 shows the result of internal erosion and piping that caused a dam failure and view of a downstream town that had to be evacuated. Since many U.S. earth dams are nearing the end of their design life (i.e., over 60 years old) [4] and facing the increasing frequency and severity of storms around the globe, it is important to find ways to efficiently monitor earth levee stability. We research the use of machine learning methods and geophysical sensor technologies to identify potential problems and better understand the structural integrity of earth levees.



Figure 1: Tunbridge Dam in Tasmania, Australia that experienced failure by internal erosion (piping) and view of evacuated town downstream (Source: Jeffery Farrar (2008) [5])

Researchers at the University of Amsterdam detected anomalies in earth levees from sensors installed inside the dams (e.g., temperature, pore water pressure, relative inclination) using a one-sided classification approach [6]. In other words, they detected deviations from what is considered a normal state of the dam or levee. Researchers at Mississippi State University experimented with unsupervised and supervised methods to detect anomalies [7] and classify levee slides [8] [9] along the Mississippi River. They investigated the use of a support vector machine to identify anomalous activity in synthetic aperture radar data. Our novel approach investigates detecting internal erosion events that could lead to failure by using geophysical data collected from sensors located on the surface of the levee, thereby retaining the integrity of the structure.

## 1.2 Detection of Anomalies

Machine learning is a branch of artificial intelligence where a computer can learn from data without human assistance. Anomaly detection is used to identify data observations that deviate from the normal or expected pattern. The detection of an anomaly in a dataset is important since an anomaly can indicate a potentially serious issue. The broad categories of anomaly detection found in the literature [10] are: supervised, semi-supervised, and unsupervised. The supervised approach requires the use of labeled training data for both normal and anomalies. This approach is often a good first step in experimentation; however, this approach is typically not the preferred method since the availability of labeled anomalous data is often unavailable or is difficult to obtain for every possible occurrence. A semi-supervised technique only requires labeled training data for what is considered normal, which is more readily available. Once a model is trained, anomalies are identified through testing the likelihood of membership. Unsupervised anomaly detection does not use labeled training data, assumes the majority class is normal, and defines outliers as anomalies in the data set. The unsupervised approach is usually the most appropriate technique for many anomaly detection problems.

The development of an anomaly detection scheme is needed in our domain of interest due to the large class imbalance of data; in other words, there is a lack of anomalous observational data in our datasets. With this approach, we can train models with what is considered normal data and detect deviations within a certain threshold. We experiment with the supervised (i.e.,

data with labels for both classes) and semi-supervised (i.e., data with labels for only the normal class). We plan to investigate unsupervised methods in future work.

## 2 Data-Driven Approach

Continuous real-time earth levee monitoring will typically produce data with the majority class representing the normal state of the structure. This normal, or baseline, data can be used as training data for our machine learning algorithms. Once trained, our detection system identifies deviations from this normal state (anomalies), which could be an indication of a developing problem or internal erosion event. We investigate techniques to address this specific one-sided classification problem, in order to build a model that can be used for the long-term health assessment of earth levees. Our approach and general data processing workflow is depicted in Figure 2 and described in the following sections.




Figure 2: General data processing and anomaly detection workflow.

### 2.1 Data Collection

We use the data set described in [1]. An experimental laboratory earth embankment was built at the United States Bureau of Reclamation (USBR) [2] and equipped with geophysical instrumentation to study internal erosion and cracking events [11]. We were provided with over 4,000 seconds of passive seismic data collected from a vertical array of geophones. The entire data set includes many crack, flow, and collapse events (30% normal or baseline data and 70% anomalous data), which is not typically representative of a live monitoring scenario. For this study, we use the first third of the dataset to more closely mimic a real-world data set (80% normal or baseline data and 20% anomalous data). Figure 3 displays the portion of the provided time-series data used in our experimentation. We continue with the data preprocessing and experiments using a support vector machine.




Figure 3: Plot of provided time-series data for a single sensor (1,380 seconds).

## 2.2 Preprocessing

**Spectral Frames** It is necessary to transform the data in order to use the dataset as input to our machine learning algorithms. The data is initially a single long-length string of time series data collected at 500Hz. The 1,380 seconds of data (which is the first third of the provided 4,140 seconds) are converted to the frequency domain, and then decomposed into segments (spectral frames) for the extraction of features. This data processing step was performed using MATLAB and the open-source MIRTtoolbox [12]. We experimented with 1, 2, 3, 5, and 10-second frame sizes and found all segmented sizes performed well (see Section 3.1 for a comparison).

**Feature Extraction** We leverage previous work for the selection of features (i.e., features commonly used in audio signal processing) to represent the internal erosion events in our passive seismic data [13]. The nine features identified and selected for our experiments are briefly described in Table 1.

Table 1: The nine spectral features extracted from each 1, 2, 3, 5, and 10-second frame [12].

Feature	Description
Zerocross (ZC)	A temporal feature which is an indicator of the noisiness of the signal. A count of the number of times the signal changes sign (crosses the zero axis).
Centroid (CR)	A statistical descriptor of spectral distribution and is the mean or geometric center (usually associated with the brightness of a sound signal).
Spread (SP)	A spectral feature that can be used to describe the asymmetry and peakedness of the signal (values represent the standard deviation of the distribution).
RMS	A temporal feature that represents the global energy (loudness) of the signal which is computed by taking the root average of the square of the amplitude.
85% Rolloff (RO)	Estimate of the amount of high frequency in the signal calculated by finding the frequency that a certain fraction (85%) of the total energy is below.
Flatness (FL)	A statistical descriptor to indicate how smooth or spiky the spectral distribution is (the ratio between the two means: geometric and arithmetic).
Kurtosis (KU)	Another statistical descriptor of the spectral distribution that can indicate the presence of transients (peaks) in the data.
Irregularity (RG)	A measurement of the extent of variation between consecutive spectral peaks.
Skewness (SK)	This feature is an indicator of the asymmetry of the spectral distribution. Zero values indicate a symmetrical distribution.

**Standardization** Using the raw feature values can impact the outcome of machine learning algorithms, especially if the feature values do not conform to a common range or variance. To address this issue in our experiments, we standardized the feature values. Data standardization is a method for transforming each of the feature vectors to have a zero mean ( $\mu$ ) and unit standard deviation ( $\sigma$ ). Standardization is the best approach for our study since we have an input space that is not sparse, are using a large margin classifier (SVM) [14], and are measuring the variance of the different features via Principal Component Analysis (PCA) [15].

**Feature Selection** Feature selection is a dimensionality reduction technique that finds a subset of the original feature set that best represents the problem space. The two most common strategies for feature selection are filter and wrapper. Filter methods are used for pre-processing and ranking of feature importance, regardless of the model selection. Wrapper methods can detect interactions between the variables and output the best performing feature subset.

### 3 Machine Learning Methods

In this paper, we experiment with two machine-learning approaches: classification and anomaly detection using the Support Vector Machine (SVM). An SVM is a supervised machine-learning algorithm primarily used for two-class classification, though an SVM can also be used to solve one-class, multi-class, and regression problems. For our classification experiment, we use the two-class SVM with all nine identified features. Given a training set of labeled data for both normal and anomalous classes, the two-class SVM aims to find a hyperplane that separates the classes while maximizing the geometric margin between them. For anomaly detection, we use a One-Class Support Vector Machine (OCSVM) with the top two ranked features identified using a feature selection algorithm (see Section 3.2). Given a training set of labeled data for only the majority class, the OCSVM seeks to find the boundary around the set of normal observations; data that lie outside of this boundary are considered the minority class, or anomalies. We apply 10-fold cross validation during our experiments and present the results next. K-fold cross validation is a technique to assess the performance of the model more accurately than a traditional single partitioning method. The idea is to partition the data into  $k$  subsets; then, for each round of validation,  $k - 1$  subsets are used for training the model and the remaining subset is used for validation. This process is repeated  $k$  times and the results are averaged.

#### 3.1 Results from Classification using a Two-Class SVM

Using an SVM with the Radial Basis Function (RBF) kernel and 10-fold cross validation, we were able to achieve 97% (or higher) overall accuracy with the nine identified features. We conducted timing tests to measure the computational cost per segment of the time series data. The entire workflow takes less than 0.0150 seconds per segment and the portion of the workflow containing the SVM runs at less than 0.0005 seconds per segment. The algorithm was trained and cross validated with the dataset described in Section 2.1. Recall that our data set contained a majority of baseline (80% normal) observations and a minority of anomalous (20% non-normal) data. We applied Principal Component Analysis (PCA) to reduce the dimensions for visualization and plot the 1st and 2nd principal components in Figure 4.




Figure 4: Data points and green contour lines show a separation of normal data from anomalies where the contour value (color bar) is zero for (a) 2-second and (b) 3-second frame sizes.

The normal data points are plotted in red, the anomalies in blue, and the support vectors are encircled in black. The color bar describes the values for our contours that provide separating boundaries between the normal observations and the anomalies based on the predicted likelihood that a label comes from a particular class. Figure 4 shows that the data is separable using an SVM and the few identified false negatives (predicted normal on anomalous data) are located outside of its appropriate group (i.e., a blue dot located in the normal data).

Using a vector of known class labels, we calculated statistics to measure the predictive performance of the algorithm and we visualize these statistics with a confusion matrix. A sample confusion matrix with descriptions is provided in Figure 5.

Output Class	0	<p><b>TP</b> True Positive</p> <ul style="list-style-type: none"> <li>• Hit</li> <li>• Predicted normal on normal data</li> </ul>	<p><b>FP</b> False Positive</p> <ul style="list-style-type: none"> <li>• False alarm</li> <li>• Predicted anomaly on normal data</li> </ul>
	1	<p><b>FN</b> False Negative</p> <ul style="list-style-type: none"> <li>• Miss</li> <li>• Predicted normal on anomalous data</li> </ul>	<p><b>TN</b> True Negative</p> <ul style="list-style-type: none"> <li>• Correct reject</li> <li>• Predicted anomaly on anomalous data</li> </ul>
		0	1
		Target Class	

Figure 5: Confusion matrix sample with descriptions.

The confusion matrices from our experiment have rows for the predicted classes and columns for the target classes. The green or diagonal elements are where the classifier predicted correctly, and the red or off-diagonal elements are where the classifier made mistakes. Figure 6a shows results from an SVM with a 2-second frame size that achieved a 98.1% overall accuracy and had only a few misses (0.7%) and false alarms (1.2%). Figure 6b displays the results from an SVM with a 3-second frame size; this larger frame size resulted in a slightly better overall accuracy of 98.9%, and also had fewer misses (0.4%) and false alarms (0.7%).




Figure 6: Confusion matrix results from our experiments with an SVM using 10-fold cross validation for (a) 2-second and (b) 3-second frame sizes.

We continued our analysis with additional metrics that are typically used in a classification problem. These metrics are values (percentages) that represent ratios of the statistics described in the confusion matrix (see Table 2).

Table 2: Evaluation metrics for machine-learning classifiers (variables are defined in Figure 5).

Metric	Description	Formula
Accuracy	All correctly predicted observations.	$\frac{TP+TN}{TotalObservations}$
Specificity	Correctly predicted anomalies.	$\frac{TN}{FP+TN}$
Precision	Correctly predicted normal class.	$\frac{TP}{TP+FP}$
Recall	Sensitivity or the hit rate.	$\frac{TP}{TP+FN}$
F1-score	Harmonic mean of precision and recall.	$2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$

Results from our runs of an SVM with 10-fold cross validation for 1, 2, 3, 5, and 10-second frame sizes are listed in Table 3. Overall, the experimental results produced over 97% accuracy and 98% F1-score (the harmonic mean between precision and recall); these results show an SVM can be used to separate normal from anomalous data observations.

Table 3: SVM 10-fold cross-validated results.

Metric	1 Second	2 Second	3 Second	5 Second	10 Second
Accuracy	0.971	0.981	0.989	0.986	0.993
Specificity	0.897	0.942	0.968	0.964	1.000
Precision	0.974	0.986	0.992	0.991	1.000
Recall	0.990	0.991	0.995	0.991	0.991
F1-score	0.982	0.988	0.993	0.991	0.995

### 3.2 Results from Anomaly Detection using a One-Class SVM

Although the nine identified spectral features performed well during our clustering [1] and two-class SVM experiments (see Section 3.1), we applied the ReliefF [16] feature selection technique to further refine the feature set for use during our one-class SVM experiments. ReliefF is a filter method for feature selection that uses k-nearest neighbors per class (we used  $k = 10$ ). Figure 7 shows the resultant weighted ranking for 1, 2, 3, 5, and 10-second frame sizes. Results indicate the top two features are zerocross and irregularity for 1 and 2-second frame sizes and zerocross and RMS for 3, 5, and 10-second frame sizes.




Figure 7: Weighted ranking from the ReliefF feature selection using our nine extracted spectral features for 1, 2, 3, 5, and 10-second frame sizes.

We used an OCSVM with the RBF kernel, 10-fold cross validation, and the top two selected features from the ReliefF algorithm. Our timing tests that measure the computational cost per segment of the time series data were slightly higher for this experiment. The entire workflow still takes less than 0.0169 seconds per segment and the portion of the workflow containing the OCSVM runs at less than 0.0007 seconds per segment. The algorithm was trained and cross validated with the same dataset described in Section 2.1. Figure 8a shows results from an OCSVM with a 2-second frame size that achieved a 83.6% overall accuracy and less than a total 16.4% misses and false alarms. Figure 8b displays the OCSVM results with a 3-second frame size which resulted in a slightly better overall accuracy of 86.5% and 13.4% total misses and false alarms.




Figure 8: Confusion matrix results from our experiments with an OCSVM using 10-fold cross validation for (a) 2-second and (b) 3-second frame sizes.

Results are listed in Table 4 for each of the 1, 2, 3, 5, and 10-second frame sizes. Even though the results are slightly reduced from our experiments with a two-class SVM, the results show the OCSVM is able to separate our normal from anomalous data observations with an accuracy of over 83% and 89% F1-score.

Table 4: OCSVM 10-fold cross-validated results.

Metric	1 Second	2 Second	3 Second	5 Second	10 Second
Accuracy	0.857	0.836	0.865	0.866	0.869
Specificity	0.641	0.600	0.667	0.714	0.750
Precision	0.909	0.898	0.916	0.926	0.934
Recall	0.913	0.896	0.916	0.905	0.900
F1-score	0.911	0.897	0.916	0.915	0.917

Figure 9 shows the results of using an OCSVM for 2-second and 3-second frame sizes. Since we used the ReliefF algorithm to reduce the dimensions, the plots represent the actual feature values for the 1st and 2nd dimensions (versus the PCA reduced data used in Section 3.1). The data points are plotted in black and the support vectors and potential outliers are encircled in red. The color bar describes the values for our contours that provide separating boundaries between the normal observations and the anomalies. As with the previous experiment with a binary SVM, the OCSVM results indicate the ability to detect anomalous data observations.




Figure 9: OCSVM results where the yellow contour shows the separation boundary around the set of normal data observations for (a) 2-second and (b) 3-second frame sizes.

Table 5 lists the mean outlier rate from the predicted class versus the actual outlier rate. The best performers are the 2-second and 3-second frame sizes with a slight decrease in performance as the frame size increases.

Table 5: OCSVM 10-fold cross-validated mean predicted versus actual outlier rates and the difference in performance (percentages). The features listed are the top two selected by the ReliefF algorithm for each specific frame size (see Table 1 for feature descriptions).

	1 Second	2 Second	3 Second	5 Second	10 Second
Predicted	20.00%	20.43%	20.21%	22.10%	23.18%
Actual	20.36%	20.29%	20.22%	20.29%	20.29%
Difference	-0.36%	0.14%	0.00%	1.81%	2.90%
Features	ZC, RG	ZC, RG	ZC, RMS	ZC, RMS	ZC, RMS

## 4 Conclusions and Future Work

Both of our experiments show that we can separate normal from anomalous data observations. We used experimental data from a laboratory earth embankment (80% normal and 20% anomalies), extracted nine features from decomposed segments of the time series data, and trained binary and one-class SVMs. We plan to continue our research for the development of a robust, generalized, automatic anomaly detection scheme that can be used to identify a developing problem or internal erosion event. To support the long-term health assessment of earth levees, our novel approach requires additional research. Further investigation of one-class learning and unsupervised techniques is required to address our specific one-sided classification problem. One-class learning is appropriate for our problem space, since it works well in the absence of anomalous data examples or in a situation where it is not possible to construct a fully representative training set of all possible non-normal observations. Unsupervised anomaly detection may also be appropriate for our domain since the model needs to perform in a continuous monitoring situation. There can also be background and spurious noise in seismic data that is caused by adjustment of equipment and machinery; earthquakes, helicopters, airplanes,

or nearby cars and trucks. We plan to research de-noising techniques (e.g., wavelet de-noising) before extracting the features. This process will help make a distinction between noise or outliers in the data and true anomalies. Finally, to create a robust, generalizable approach, our workflow must be tested on different types of earth dam and levee passive seismic data sets. Additional data that is available for our experimentation include: IJkdijk full-scale test embankment and Colijnsplaat real-world levee (both located in the Netherlands).

**Acknowledgments.** This work is supported in part by National Science Foundation Grant OISE-1243539.

## References

- [1] W. Belcher, T. Camp, and V. V. Krzhizhanovskaya, “Detecting erosion events in earth dam and levee passive seismic data with clustering,” *Proceedings of the 14th International Conference on Machine Learning and Applications (ICMLA)*, pp. 903–910, 2015.
- [2] “US Bureau of Reclamation,” <http://www.usbr.gov>, accessed: 2015-10-23.
- [3] M. Foster, R. Fell, and M. Spannagle, “The statistics of embankment dam failures and accidents,” *Canadian Geotechnical Journal*, vol. 37, no. 5, pp. 1000–1024, 2000.
- [4] “Aging water resource infrastructure in the United States,” <http://www.usbr.gov/newsroom/testimony/detail.cfm?RecordID=2441>, accessed: 2015-10-23.
- [5] “Piping and internal erosion failure, Tunbridge Dam, Tasmania, Australia,” <http://www.geoengineer.org/gallery/Earthfill+-+Rockfill+Dams/Tunbridge+Dam/>, accessed: 2015-10-23.
- [6] A. L. Pyayt, A. P. Kozionov, I. I. Mokhov, B. Lang, R. J. Meijer, V. V. Krzhizhanovskaya, and P. M. Sloot, “Time-frequency methods for structural health monitoring,” *Sensors*, vol. 14, no. 3, pp. 5147–5173, 2014.
- [7] L. Dabbiru, J. V. Aanstoos, M. Mahrooghy, W. Li, A. Shanker, and N. H. Younan, “Levee anomaly detection using polarimetric synthetic aperture radar data,” *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 5113–5116, 2012.
- [8] L. Dabbiru, J. V. Aanstoos, and N. H. Younan, “Earthen levee slide detection via automated analysis of synthetic aperture radar imagery,” *Landslides*, pp. 1–10, 2015.
- [9] D. Han, Q. Du, J. V. Aanstoos, and N. Younan, “Classification of levee slides from airborne synthetic aperture radar images with efficient spatial feature extraction,” *Journal of Applied Remote Sensing*, vol. 9, no. 1, pp. 097294–1–097294–10, 2015.
- [10] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Computing Surveys (CSUR)*, vol. 41, no. 3, pp. 1–72, 2009.
- [11] R. V. Rinehart, M. L. Parekh, J. B. Rittgers, M. A. Mooney, and A. Revil, “Preliminary implementation of geophysical techniques to monitor embankment dam filter cracking at the laboratory scale,” *Proceedings of the 6th Annual International Conference on Software Engineering (ICSE)*, 2012.
- [12] O. Lartillot and P. Toivainen, “A Matlab toolbox for musical feature extraction from audio,” *Proceedings of the 10th International Conference on Digital Audio Effects*, 2007.
- [13] M. J. Rubin, “Efficient and automatic wireless geohazard monitoring,” Ph.D. dissertation, Colorado School of Mines, 2014.
- [14] A. Ben-Hur and J. Weston, “A user’s guide to support vector machines,” *Data Mining Techniques for the Life Sciences*, pp. 223–239, 2010.
- [15] I. Jolliffe, *Principal component analysis*. Wiley Online Library, 2002.
- [16] M. Robnik-Šikonja and I. Kononenko, “Theoretical and empirical analysis of ReliefF and RReliefF,” *Machine learning*, vol. 53, no. 1-2, pp. 23–69, 2003.