

Detecting Seismic Events with Computer Vision: Applications for Fiber-Optic Sensing

Solvi Thrastarson¹, Robert Torfason², Sara Klaasen³, Patrick Paitz³, Yesim CUBUK SABUNCU⁴, Kristín Jónsdóttir⁵, and Andreas Fichtner⁶

¹ETH Zurich, Switzerland

²Genki Instruments

³Swiss Federal Institute of Technology in Zurich

⁴Icelandic Meteorological Office


⁵Icelandic Met Office, Iceland


⁶ETH Zurich


November 23, 2022

Abstract

Seismologists working with fiber-optic sensing, commonly referred to as Distributed Acoustic Sensing (DAS), have yet to find an established way of automatically detecting signals of interest within its recordings. We propose a new research perspective within the field by examining the output of a DAS array as an image and processing the image to find signals of interest. In this manuscript, we show an example of such a method, where we automatically detect seismic events of interest within two different DAS datasets, finding, respectively 99 % and 96 % of the local earthquakes previously identified within the data by manual analysis. The method is based on simple image processing and computer vision techniques, which clean the image, and, ideally, leave nothing but the signal of interest. These simple image processing steps yield promising results, indicating that computer vision and image processing might have an immediate impact in geophysical applications of fiber-optic sensing.







Detecting Seismic Events with Computer Vision: Applications for Fiber-Optic Sensing

Solvi Thrastarson¹, Robert Torfason², Sara Klaasen¹, Patrick Paitz¹, Yeşim
Çubuk-Sabuncu³, Kristín Jónsdóttir³, Andreas Fichtner¹

¹ETH Zürich, Sonneggstrasse 5, 8092, Zürich, Switzerland

²Genki Instruments, Klapparstígur 25-27, 101 Reykjavík, Iceland

³Icelandic Meteorological Office, Bústaðavegur 9, 150 Reykjavík, Iceland

Key Points:

- DAS datasets are too large for manual inspection, and detection algorithms built for seismometers do not exploit the benefits of DAS.
- Recordings from DAS arrays can be viewed as images, making it possible to use computer vision techniques to detect earthquakes.
- We develop a computer vision pipeline that automatically detects seismic events and we test it on two separate datasets.

Abstract

Seismologists working with fiber-optic sensing, commonly referred to as Distributed Acoustic Sensing (DAS), have yet to find an established way of automatically detecting signals of interest within its recordings. We propose a new research perspective within the field by examining the output of a DAS array as an image and processing the image to find signals of interest. In this manuscript, we show an example of such a method, where we automatically detect seismic events of interest within two different DAS datasets, finding, respectively 99% and 96% of the local earthquakes previously identified within the data by manual analysis. The method is based on simple image processing and computer vision techniques, which clean the image, and, ideally, leave nothing but the signal of interest. These simple image processing steps yield promising results, indicating that computer vision and image processing might have an immediate impact in geophysical applications of fiber-optic sensing.

Plain Language Summary

Detecting earthquakes in large volumes of data becomes challenging, as the data can no longer be manually inspected within a reasonable timeframe. We formulate a detection algorithm that automatically detects the presence of earthquakes in a large data volume. We treat the data as an image, making the earthquakes visually distinctive from the noise. Our method is based on image processing techniques, and, step-by-step, we process the data until only physical signals of interest remain in the image. The algorithm has been tested on two different data sets, finding 99% and 96%, respectively, of all local earthquakes. This is a promising result, showing the potential of image-based processing techniques to deal with sizeable seismic data volumes.

1 Introduction

Distributed Acoustic Sensing (DAS) is an emerging technology in geophysics that yields a densely sampled array of seismic measurements. An interrogation unit sends light pulses into a fiber-optic cable, and naturally occurring heterogeneities in the cable cause back-scattering of the light pulses (Hartog, 2017). The axial strain of the cable causes phase shifts in the back-scattered signals as it deforms. The entire fiber-optic cable is sampled, and the recorded phase shifts are averaged over a gauge length, generally a few meters. The result is a dense seismic network with closely-spaced stations (down to 25 cm) that extends up to several kilometers with a high temporal sampling rate (up to several kHz).

The high spatial and temporal resolutions of DAS create large data volumes (up to tens of TB) that can be challenging to process and store. An effort is made to automate processing procedures of DAS measurements, as manual processing becomes infeasible. However, automated processing techniques that have been developed for seismometers do not necessarily translate well to data collected with DAS for the following reasons. (1) They are designed for either three-component or single vertical component sensors instead of uni-directional strain measurements. (2) They do not harness the dense spatial sampling which may enable the identification of additional low magnitude events. (3) Applying them individually to each DAS channel might become prohibitively expensive computationally.

1.1 Event Detection Algorithms

Many earthquake detection algorithms have been developed for single seismic stations throughout the years, e.g., short-term average over long-term average (Allen, 1978). Among the most established detection algorithms are similarity-based algorithms, such as template matching (e.g. Shelly et al., 2007; Peng & Zhao, 2009; Skoumal et al., 2015).

Using cross-correlation, previously recorded waveforms (templates) from seismically active regions are used to detect comparable earthquakes in similar areas. Machine learning methods have been developed to detect events (e.g. Perol et al., 2018; Ross et al., 2018), classify them (e.g. Falsaperla et al., 1996; Masotti et al., 2006; Maggi et al., 2017), or both (e.g. Beyreuther & Wassermann, 2008; Hammer et al., 2013; Dammeier et al., 2016; Wenner et al., 2021).

DAS poses a challenge for template matching as it is often deployed as a temporary array where a catalog of recurring waveforms does not exist. To employ template matching with a temporary DAS array, a careful selection of template waveforms needs to be done a priori. Given the data volumes recorded by a DAS array, searching for a collection of templates may already be prohibitive. To circumvent this problem, Li and Zhan (2018) used a catalog of events detected by a permanent seismic array to find template waveforms recorded by the DAS array. They showed that this could help detect events with a low signal-to-noise ratio. It does, however, not guarantee that the template events represent all the signals of interest that the array might record, and the computational requirements may be high due to the need for two-dimensional correlations.

The use of single channels of the DAS array to detect earthquakes does not fully exploit the potential of a DAS array, which is its dense, spatial sampling. Consequently, developments within the field of array seismology may be better suited to exploit the high spatial resolution of DAS. Chmiel et al. (2019) developed a matched field processing method that detects and localizes events with seismic arrays. Such methods could translate well to DAS measurements. However, they would require significant adaptations to take the in-line strain-rate measurements into account, and the fiber-optic cable layout would need to be planned with such a method in mind.

Rather than considering individual DAS measurements, we propose to visualize the data as a two-dimensional image where the axes represent space and time, and the color shows the intensity of the recorded signals. Signals that are challenging to identify on individual channels become apparent when visualizing the entire array. We aim to exploit the patterns in the data caused by earthquakes in order to detect them. We use computer vision techniques to identify the patterns in two-dimensional images, making it a potential earthquake detection tool.

1.2 Computer Vision and Image Processing

On a high level, computer vision is image processing with the goal of image understanding. The image processing involves algorithms to reduce noise or enhance certain features of an image (Sonka et al., 2014). Image processing is usually a step within computer vision pipelines prior to the final step of image understanding. As an example, for a typical neural network, the values of the image need to be normalized to a specific range and resized before the image can be interpreted (He et al., 2016).

Most of the recent advances in computer vision have used convolutional neural networks, and typical tasks for modern computer vision systems include classification (e.g. He et al., 2016), object detection (e.g. Redmon & Farhadi, 2017) and image segmentation (e.g. L. Chen et al., 2017). There are, however, drawbacks related to using deep learning methods (neural networks) in the small-data regime with a heavy class imbalance. Large volumes of data are needed to train those neural networks, and the data need to be meticulously labeled, which is time-consuming. Without a sufficient amount of data, the task could be impossible to solve with deep learning methods (Karimi et al., 2020).

While a single DAS survey does not necessarily create enough data to train a neural network properly, the quantities of data are still large enough to make it prohibitively time-consuming for manual inspection and expensive for long-term storage.

A large majority of passive recordings do not include any signals of interest, and it is thus a viable option to only store the recordings which contain interesting signals and delete the rest. As the datasets are too large to inspect manually, we need a method that automatically detects seismic events within the dataset. Ideally, we do not want to

miss any interesting data, so the occurrence of false-negative results needs to be minimized.

Recent applications of machine learning techniques with DAS and conventional seismic data focused on de-noising of data (Beckouche & Ma, 2014; Martin et al., 2018; Zhu et al., 2019; Yu et al., 2019; Y. Chen et al., 2019; van den Ende et al., 2021) and arrival-time picking (Zhu & Beroza, 2019). There have been attempts at using deep learning based computer vision techniques to detect seismic events within DAS recordings. Binder and Chakraborty (2019) and Stork et al. (2020) applied a convolutional neural network to the measurements from a downhole DAS system to detect microseismic events. They trained the network using synthetic data superimposed with real recorded ambient noise.

Whereas these detection methods show promise on synthetic data, they have a problem producing robust results with actual data. Even though the neural network detected multiple events in both cases, it also missed many events, which we aim to minimize with our proposed detection algorithm.

Even though non-deep learning methods may be less, they are usually easy to understand and can be tuned based on prior information about the task (Otsu, 1979; Duda & Hart, 1972; Vincent, 1993). They can thus be a logical first step before using deep learning based methods.

Neural networks are especially useful in object detection when the characteristics of the objects cannot be described to a computer easily (e.g., a cat) (Girshick et al., 2014). However, when the general characteristics of the objects can be described, more straightforward approaches become more accessible and are easier to tune.

The approach we present here is, to our knowledge, the first step in the direction of event detection from DAS recordings using simple image processing techniques. We believe that the method still has unfulfilled potential, yet we want to describe our current solution and encourage people to further experiment with such an approach.

2 Description of Algorithm

We designed the algorithm to identify a signal of interest from the background noise with image processing techniques. To achieve this, it is essential to differentiate the signal from the noise. This includes ambient and anthropogenic noise, as well as instrumental noise, which may appear in the form of horizontal stripes, as can be seen in Figure 1.

The data are initially represented as a two-dimensional array, where the recorded strain-rate is represented in space and time. The strain-rate is color-coded, yielding a two-dimensional image, such as the top row of Figure 1. The image processing algorithm works on inverted grayscale images, while the top row of Figure 1 is represented in intensities of red, green, and blue (*RGB*). We convert from *RGB* to inverted grayscale (*G*) with a formula from Poynton (1997),

$$G = 1.0 - \frac{0.2125R + 0.7154G + 0.0721B}{255}, \quad (1)$$

which results in an image where brightness is a proxy for strain-rate intensity. The image can be viewed as a two-dimensional array of strain-rate intensities where time is on the x-axis, and DAS channels are on the y-axis. Using our prior knowledge of how a seismic event would distinguish itself from the noise, we constructed an algorithm that reduces the noise step-by-step, ideally leaving only the signal of interest in the image. The steps are described in detail in the following subsections (along with their tuning parameters), and a visual guide is provided in Figure 1 along with references to relevant section numbers.

2.1 General Noise Reduction

The noise level of a DAS measurement often varies between channels as they are separated in space, and coupling to the ground varies. To lower the noise level of the image, we remove the median and/or mean brightness per channel, bringing anything be-




Figure 1. The algorithm shown step-by-step for three different cases. Each step is mentioned above the row of figures where the step has been applied, and the numbers at the right end of the columns refer to the section numbers in the text where the step is explained. The grayscale colormap is inverted for easier visibility. From the thresholding step, black denotes True, and white denotes False. The steps are detailed in section 2. **a)**: Small earthquake in top left corner detected. Two True regions left after processing. **b)**: No data of interest, nothing detected. No True regions left after processing. **c)**: Large earthquake on the right side detected. 60 True regions left after processing. All figures have the same axes, 30 s time window on the horizontal one and 12 km along cable on the vertical one.

low the median/mean intensity of its respective channel to zero. The tuning possibilities here are whether to remove the median, the mean, or both, and which channels to include in the image. Parts of the fiber-optic cable may have such bad coupling to the ground that it affects the channels around it. This happens, for example, on a glacier when a part of the cable is hanging in the air between the two sides of a crevasse. In such a case, when that part of the cable vibrates due to wind, the vibrations can propagate along the cable, which would cause multiple false detections. One way to combat such an effect is to remove the relevant channels from the image. Preparing the data for the detection algorithm is application-specific and will be further discussed in section 4.

2.2 Brightness Thresholding

Many image processing algorithms work on binary images. Therefore, we apply a thresholding algorithm to convert the data into binary images. The thresholding decides at which intensity level the signal is classified as True, and below it, as False.

The method we use for thresholding is Otsu's method (Otsu, 1979). The method can be applied to find a single threshold to separate images into foreground/background. It can also be used to find several thresholds to further separate the image into different classes. The thresholding uses the histogram of pixel brightness and separates it into two classes by minimizing the intra-class variance, or equivalently maximizing the inter-class variance. This can also be seen as maximizing the distance between the average brightnesses of the clusters. The tuning parameters here are the number of thresholds and at which values the True/False boundary lies.

2.3 Coherency Thresholding

The binary image will include multiple True regions, most of which not being coherent in space. This results from the noise in the recordings not being coherent between channels. Any actual signal of interest will consist of multiple contiguous True pixels, as it is a physical signal recorded by multiple channels. As we consider earthquake signals, the signals have to satisfy the wave equation. A signal which satisfies the wave equation will propagate along the array, displaying a coherent signal from its onset at an angle between 0° and 90° . In contrast, instrumental noise appears as thin, horizontal stripes, as can be seen in Figure 1. That is why we can clean the image by setting every True pixel to False if it is not a part of a larger cluster of True pixels. The tuning parameter here is the minimum size that a cluster of True pixels needs to have to remain True.

2.4 Remove Unwanted Shapes

As previously mentioned, the signal of interest propagates along the array in space and time and is thus neither strictly horizontal nor vertical within an image. We can thus safely remove any True pixels which are a part of a narrow horizontal or vertical line of Trues. This is done by making a horizontal/vertical template move through the image. When the template is matched, purely horizontal/vertical patterns are deleted. When a signal arrives from directly below the array, it can affect the whole array simultaneously, resulting in a vertical line in the image. The vertical line removal does not affect such signals, as it will remove only narrow vertical stripes, and a broader pattern in the image due to an earthquake will thus not be removed.

This step is not limited to horizontal/vertical lines. The template can be any expected, unwanted signal e.g., a car driving along a cable in urban seismology. In the presented examples, we will limit ourselves to horizontal and vertical lines as the measurements were taken in remote environments. The tuning parameter here is the size and shape of the templates.

2.5 Final Cleaning of Image

Removing the horizontal and vertical lines does not always work consistently, as the width of the lines varies, and it occurs that lines are only partially removed. However, the small noisy lines get interrupted, which reduces their True cluster size. A repetition of the same method as described in section 2.3 will now remove the smaller True clusters that remain of the horizontal and vertical lines. After this step, virtually all noise has been removed from the image. Therefore, the proposed algorithm can be described as an image cleaning process, where the remaining signal is likely to be of interest.

2.6 Count Regions

As the remaining signal is likely to be some coherent physical signal (i.e., earthquake, volcanic tremor, icequake, explosion), we can count the number and size of True regions to estimate a probability of an event of interest being the signal's origin. This is a tuning parameter that can affect the true/false positives/negatives of the results. The optimal value of this parameter depends on the objective of the algorithm. In the example below, we assigned the earthquake label to any image which was not fully False, as our objectives were primarily to minimize the loss of any physical signals, and secondly, to reduce the data volume.

3 Test Case

We tested the classifier on two existing datasets, described in details below. To enable the establishment of a ground truth catalog, we limited ourselves to subsets of the two datasets that are small enough to permit manual analysis. The length of our time window in the examples is 30 s, as we expect it to be large enough to contain the signal of an entire event and small enough to make it unlikely to have more than one event present. Our presented algorithm is not limited to 30 s windows and can scale to bigger or smaller window sizes.

The application that this classifier was initially designed for is to reduce the amount of data from a DAS survey without losing any signal of interest. The goal is thus twofold: to make sure no data of interest gets lost and reduce the amount of data that needs closer attention and further analysis as much as possible. Ideally, we would like to delete the data that the classifier marks as uninteresting for our purposes.

For both tested datasets, the algorithm was tuned by taking 3-4 examples of different types of signals and noise from different times in the dataset. The algorithm was tuned to classify these correctly and then run with these settings on the entire test periods. The example signals can be events of various sizes plus one or more examples without a signal of interest. Increasing diversity in the signals used for tuning, results in increased generalization to unseen measurements.

When evaluating the performance of a binary classifier with a high class-imbalance such as the one presented in this study, two metrics are generally used. Firstly, it is *precision*,

$$PR = \frac{TP}{TP + FP}, \quad (2)$$

where TP and FP are the numbers of true and false positives, respectively. Secondly, it is *recall*,

$$RC = \frac{TP}{TP + FN}, \quad (3)$$

where FN stands for the number of signals falsely classified as negative. Precision measures how many of the detections were correct, while recall measures the portion of the signals of interest which were detected. There is a trade-off between the two metrics, so classifiers are often tuned to optimize more for one while trying to achieve reasonable performance in the other (Davis & Goadrich, 2006).




Figure 2. The cable layout of the Grímsvötn dataset (left) and Mount Meager dataset (right). The cables are represented with a red line.

With our primary objective of not losing any signal of interest, we optimized more for recall rather than precision.

3.1 Grímsvötn

Grímsvötn is Iceland’s most active volcano, with its most recent eruption occurring in 2011. It is situated underneath the Vatnajökull ice cap, and our DAS dataset was acquired by trenching a fiber-optic cable 50 cm into the snow cover. The total length of the cable was 12 km, and its layout can be seen in Figure 2. A Silixa iDAS(TM) interrogation unit was running for one month from May to June 2021, generating approximately one terabyte of data. Due to the deep trenching, the dataset has a relatively high signal-to-noise ratio. Physical signals were easily visible within the data by looking at images such as the top right one in Figure 1, which inspired the construction of the presented detection algorithm. The performance evaluation was done for a single day of data, as it is still at a size where the creation of ground truth is possible within a reasonable amount of time. The data were converted from light intensity to nanostrain-rate, and no further processing was applied. The results of the evaluation can be seen in Table 1, where it is evident that out of the 145 signals of interest manually detected during the day, 143 were detected with our proposed detection algorithm. The misclassified seismic events both contained weak signals barely visible by eye. By running our new detection algorithm on the dataset, 99 % of the signals, i.e., 143 out of 145 events, previously defined as signals of interest, were recovered. Out of the total 2880 time windows, 196 were classified to contain a signal of interest, meaning that only 7 % of the original data require longer-term storage.

3.2 Mount Meager

Mount Meager is an active volcano in the Garibaldi Volcanic Belt in British Columbia, Canada. The area is characterized by a range of natural hazards, as it was home to the largest recorded landslide in Canada in 2010 (Read, 1990; Roberti et al., 2018). In addition, the environment may be further destabilized by the melting glaciers due to climate change. However, the area also has the greatest geothermal potential of Canada (Jessop et al., 1991). The DAS experiment on Mount Meager was part of a larger project to map the feasibility of geothermal exploration on the Mount Meager complex.

A 3 km long fiber-optic cable was installed along a ridge of Mount Meager and interrogated by an OptaSense ODH3 interrogation unit continuously for one month between September and October 2019 (Klaasen et al., 2021). The cable layout can be seen in Figure 2.

We evaluated the classification performance of our algorithm over two days of data. The performance is displayed in Table 1. The data were converted from light intensity to nanostrain-rate and bandpass filtered between 5 and 45 Hz. The algorithm found 96 % of the events previously identified during these two days, within only 4 % of the entire dataset, again easing the inspection of the signals of interest. The identified events which were not detected by our algorithm had signals which were too faint to pass the brightness thresholding step.

Table 1. Classification table for the test cases. The algorithm achieves 99 % recall and 73 % precision for the Grímsvötn test case. It achieves 96 % recall and 60% precision for Mount Meager

		Classification			
		Grímsvötn		Mount Meager	
		Positive	Negative	Positive	Negative
Truth	Positive	143	2	104	4
	Negative	53	2682	69	4672

4 Discussions and Conclusions

We have presented an example of an algorithm that can automatically detect seismic events recorded with a DAS array. The algorithm is simplistic and easy to understand while still showing promising results on two different datasets. This hints at the potential of such an approach, and we aim to take the idea further in future work.

An emerging problem with DAS is that the data volumes are simply too large to, on the one hand, analyze manually, and on the other hand, store. An application like this enables working with a small fraction of the dataset, making it easier to both process and store event-specific data in the long term.

So far, we have used the method to create an event catalog for Grímsvötn by running the autodetection on the entire dataset and then manually going through the event detections to find false positives. With more development, we are confident that the manual part of that workflow can be cut down further or even removed entirely.

Preprocessing of the waveforms can greatly affect the performance of the algorithm, and we believe that it could improve the performance of the classifier. One such method is frequency-wavenumber filtering, which can be applied to DAS recordings due to its high spatial and temporal resolution. This could further remove unwanted noise in the images, easing the image processing afterward.

After a number of seismic events have been identified, a diverse selection of events can be isolated as two-dimensional templates. These templates could be used in two-dimensional cross-correlation to look for similar earthquakes within the data, creating a DAS variant of the template matching algorithm.

Currently, the presented tool works as a reliable way of reducing the data volume to work with while not losing a significant part of the signals of interest. We believe that with further work in this direction, the algorithm's precision can be significantly increased, making it a reliable way to automatically create earthquake catalogs from DAS arrays.

In order to exploit the power of deep learning based methods to the fullest, the results from DAS experiments need to be generalized in a way that a single neural network can read the results of various experiments. By doing that, enough data could be gath-

ered to train a fully convolutional neural network to detect earthquakes in images from various DAS surveys.

5 Code and Data Availability Statement

The dataset collected on Grímsvötn is currently not publicly available. However, the images used to test the algorithm are available (Thrastarson et al., 2021). The dataset collected on Mount Meager is archived at BC Geoscience, and is available upon request. The code to use the detection algorithm (Thrastarson, 2021) is publicly open and available along with a showcase Binder notebook: www.github.com/solvithrastar/DAS_Auto.

Acknowledgments

For the processing of the data, we used Python (Rossum, 1995). Specifically NumPy (Harris et al., 2020), Matplotlib (Hunter, 2007), SciPy (Virtanen et al., 2020), ObsPy (Beyreuther et al., 2010; Krischer et al., 2015), Scikit-image (Van der Walt et al., 2014) and OpenCV (Bradski & Kaehler, 2008).

This work was supported by the European Unions Horizon 2020 research and innovation program through an ERC Starting Grant (The Collaborative Seismic Earth Model, grant No. 714069). The real-time Earthquake Risk Reduction for a Resilient Europe 'RISE' project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 82111. We acknowledge the help of people who assisted in acquiring the datasets on which the algorithm was tested.

For Grímsvötn: Bergur H. Bergsson, Bergur Einarsson, Guðlaugur Jakob Þorsteinson, Hildur Jónsdóttir, Hlynur Skagfjörð Pálsson, Jóhannes Rögnvaldsson, Laufey Guðmundsdóttir, Pálmi Erlendsson, Snæbjörn Sveinsson, Vala Hjörleifsdóttir, Vilhjálmur Kjartansson, Neyðarlínan and The Icelandic Glaciological Society as well as Athena Chalari and the Silixa support team.

For Mount Meager: Jan Dettmer, Hersh Gilbert, Jackie Smale, and Optasense. The Mount Meager data acquisition was funded by the Government of Canada's New Frontiers in Research Fund (NFRF).

References

- Allen, R. V. (1978). Automatic earthquake recognition and timing from single traces. *Bull. Seismol. Soc. Am.*, 68(5), 1521–1532.
- Beckouche, S., & Ma, J. (2014). Simultaneous dictionary learning and denoising for seismic data. *Geophysics*, 79(3), A27–A31.
- Beyreuther, M., Barsch, R., Krischer, L., & Wassermann, J. (2010). ObsPy: A Python toolbox for seismology. *Seis. Res. Lett.*, 81, 47–58.
- Beyreuther, M., & Wassermann, J. (2008). Continuous earthquake detection and classification using discrete hidden markov models. *Geophys. J. Int.*, 175(3), 1055–1066.
- Binder, G., & Chakraborty, D. (2019). Detecting microseismic events in downhole distributed acoustic sensing data using convolutional neural networks. In *Seg technical program expanded abstracts 2019* (pp. 4864–4868). Society of Exploration Geophysicists.
- Bradski, G., & Kaehler, A. (2008). *Learning opencv: Computer vision with the opencv library*. " O'Reilly Media, Inc."
- Chen, L., Papandreou, G., Schroff, F., & Adam, H. (2017). Rethinking atrous convolution for semantic image segmentation. *CoRR*, abs/1706.05587. Retrieved from <http://arxiv.org/abs/1706.05587>
- Chen, Y., Zhang, M., Bai, M., & Chen, W. (2019). Improving the signal-to-noise ratio of seismological datasets by unsupervised machine learning. *Seismological Research Letters*, 90(4), 1552–1564.

- Chmiel, M., Roux, P., & Bardainne, T. (2019). High-sensitivity microseismic monitoring: Automatic detection and localization of subsurface noise sources using matched-field processing and dense patch arrays. *Geophysics*, 84(6), KS211–KS223.
- Dammeier, F., Moore, J. R., Hammer, C., Haslinger, F., & Loew, S. (2016). Automatic detection of alpine rockslides in continuous seismic data using hidden markov models. *J. Geophys. Res.*, 121(2), 351–371. doi: <https://doi.org/10.1002/2015JF003647>
- Davis, J., & Goadrich, M. (2006). The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on machine learning* (p. 233–240). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/1143844.1143874> doi: 10.1145/1143844.1143874
- Duda, R. O., & Hart, P. E. (1972, January). Use of the hough transformation to detect lines and curves in pictures. *Commun. ACM*, 15(1), 11–15. Retrieved from <https://doi.org/10.1145/361237.361242> doi: 10.1145/361237.361242
- Falsaperla, S., Graziani, S., Nunnari, G., & Spampinato, S. (1996). Automatic classification of volcanic earthquakes by using multi-layered neural networks. *Natural Hazards*, 13(3), 205–228.
- Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *2014 ieee conference on computer vision and pattern recognition* (p. 580–587). doi: 10.1109/CVPR.2014.81
- Hammer, C., Ohrnberger, M., & Fäh, D. (2013). Classifying seismic waveforms from scratch: a case study in the alpine environment. *Geophys. J. Int.*, 192(1), 425–439.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., ... others (2020). Array programming with numpy. *Nature*, 585(7825), 357–362.
- Hartog, A. H. (2017). *An introduction to distributed optical fibre sensors*. CRC press.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *2016 ieee conference on computer vision and pattern recognition (cvpr)* (p. 770–778). doi: 10.1109/CVPR.2016.90
- Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(03), 90–95.
- Jessop, A. M., Ghomshei, M. M., & Drury, M. J. (1991). Geothermal energy in Canada. *Geothermics*, 20(5-6), 369–385.
- Karimi, D., Dou, H., Warfield, S. K., & Gholipour, A. (2020). Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis. *Medical Image Analysis*, 65, 101759. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1361841520301237> doi: <https://doi.org/10.1016/j.media.2020.101759>
- Klaasen, S., Paitz, P., Lindner, N., Dettmer, J., & Fichtner, A. (2021). Distributed acoustic sensing in volcano-glacial environments—mount meager, british columbia. *J. Geophys. Res.*, 126(11), e2021JB022358. doi: 10.1029/2021JB022358
- Krischer, L., Megies, T., Barsch, R., Beyreuther, M., Lecocq, T., Caudron, C., & Wassermann, J. (2015). ObsPy: A bridge for seismology into the scientific Python ecosystem. *Computational Science & Discovery*, 8(1), 014003. doi: 10.1088/1749-4699/8/1/014003
- Li, Z., & Zhan, Z. (2018). Pushing the limit of earthquake detection with distributed acoustic sensing and template matching: A case study at the brady geothermal field. *Geophys. J. Int.*, 215(3), 1583–1593.

- Maggi, A., Ferrazzini, V., Hibert, C., Beauducel, F., Boissier, P., & Amemoutou, A. (2017). Implementation of a multistation approach for automated event classification at piton de la fournaise volcano. *Seism. Res. Lett.*, 88(3), 878–891. doi: <https://doi.org/10.1785/0220160189>
- Martin, E. R., Huot, F., Ma, Y., Cieplicki, R., Cole, S., Karrenbach, M., & Biondi, B. L. (2018). A seismic shift in scalable acquisition demands new processing: Fiber-optic seismic signal retrieval in urban areas with unsupervised learning for coherent noise removal. *IEEE Signal Processing Magazine*, 35(2), 31–40.
- Masotti, M., Falsaperla, S., Langer, H., Spampinato, S., & Campanini, R. (2006). Application of support vector machine to the classification of volcanic tremor at etna, italy. *Geophys. Res. Lett.*, 33(20).
- Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1), 62–66.
- Peng, Z., & Zhao, P. (2009). Migration of early aftershocks following the 2004 park-field earthquake. *Nature Geoscience*, 2(12), 877–881.
- Perol, T., Gharbi, M., & Denolle, M. (2018). Convolutional neural network for earthquake detection and location. *Science Advances*, 4(2), e1700578.
- Poynton, C. (1997). Frequently asked questions about color. Retrieved June, 19, 2004.
- Read, P. B. (1990). Mount Meager Complex, Garibaldi Belt, southwestern British Columbia. *Geoscience Canada*.
- Redmon, J., & Farhadi, A. (2017). Yolo9000: Better, faster, stronger. In *2017 IEEE conference on computer vision and pattern recognition (cvpr)* (p. 6517-6525). doi: 10.1109/CVPR.2017.690
- Roberti, G., Ward, B., van Wyk De Vries, B., Falorni, G., Menounos, B., Friele, P., ... others (2018). Landslides and glacier retreat at Mt. Meager volcano: Hazard and risk challenges. *Geohazards Engineering Resiliency in a Changing Climate*.
- Ross, Z. E., Meier, M.-A., Hauksson, E., & Heaton, T. H. (2018). Generalized seismic phase detection with deep learning. *Bull. Seismol. Soc. Am.*, 108(5A), 2894–2901.
- Rossum, G. (1995). Python reference manual.
- Shelly, D. R., Beroza, G. C., & Ide, S. (2007). Non-volcanic tremor and low-frequency earthquake swarms. *Nature*, 446(7133), 305–307.
- Skoumal, R. J., Brudzinski, M. R., & Currie, B. S. (2015). Earthquakes induced by hydraulic fracturing in poland township, ohio. *Bull. Seismol. Soc. Am.*, 105(1), 189–197.
- Sonka, M., Hlavac, V., & Boyle, R. (2014). *Image processing, analysis, and machine vision*. Cengage Learning.
- Stork, A. L., Baird, A. F., Horne, S. A., Naldrett, G., Lapins, S., Kendall, J.-M., ... Williams, A. (2020). Application of machine learning to microseismic event detection in distributed acoustic sensing data. *Geophysics*, 85(5), KS149–KS160.
- Thrastarson, S. (2021, December). *solvithraster/DAS_Auto: DAS_Auto - Image based earthquake detection algorithms*. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.5768159> doi: 10.5281/zenodo.5768159
- Thrastarson, S., Klaasen, S., Cubuk-Sabunku, Y., Jónsdóttir, K., & Fichtner, A. (2021, December). *Images of das recordings from grímsvötn*. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.5769827> doi: 10.5281/zenodo.5769827
- van den Ende, M., Lior, I., Ampuero, J.-P., Sladen, A., Ferrari, A., & Richard, C. (2021). A self-supervised deep learning approach for blind denoising and waveform coherence enhancement in distributed acoustic sensing data.
- Van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., ... Yu, T. (2014). scikit-image: image processing in python. *PeerJ*, 2, e453.

487 Vincent, L. (1993). Grayscale area opening and closing, their efficient implementa-
488 tion and applications. In *in proc. eurasip workshop on mathematical morphol-*
489 *ogy and its applications to signal processing.*


490 Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Courn-
491 peau, D., ... others (2020). Scipy 1.0: fundamental algorithms for scientific
492 computing in python. *Nature methods*, 17(3), 261–272.

493 Wenner, M., Hibert, C., van Herwijnen, A., Meier, L., & Walter, F. (2021). Near-
494 real-time automated classification of seismic signals of slope failures with
495 continuous random forests. *Natural Hazards and Earth System Sciences*, 21(1),
496 339–361. doi: <https://doi.org/10.5194/nhess-21-339-2021>


497 Yu, S., Ma, J., & Wang, W. (2019). Deep learning for denoising. *Geophysics*, 84(6),
498 V333–V350.

499 Zhu, W., & Beroza, G. C. (2019). Phasenet: a deep-neural-network-based seismic
500 arrival-time picking method. *Geophys. J. Int.*, 216(1), 261–273.


501 Zhu, W., Mousavi, S. M., & Beroza, G. C. (2019). Seismic signal denoising and
502 decomposition using deep neural networks. *IEEE Transactions on Geoscience*
503 *and Remote Sensing*, 57(11), 9476–9488.



RGB converted to grayscale



Median/Mean intensity removed per channel



2.1

Brightness thresholding



2.2

Coherency thresholding



2.3

Horizontal and Vertical Trues removed



2.4

Coherency thresholding



2.5

Distance (12 km)

Time (30 s)

a)

b)

c)

HUT



12 km



