

HUMAN AND MACHINE LISTENING OF SEISMIC DATA

Arthur Paté, Benjamin Holtzman, Felix Waldhauser, Douglas Repetto John Paisley

Lamont-Doherty Earth observatory
Columbia University
Palisades, NY, USA
pate@ldeo.columbia.edu

Department of Electrical Engineering
Data Science Institute
Columbia University
New York, NY, USA

ABSTRACT

Geothermal energy mining consists of injecting cold water into hot rocks in order to create micro-fractures allowing heat to be extracted and converted into electrical energy. This water injection can trigger several rock fracture processes. Seismologists are facing the challenge of identifying and understanding these fracture processes in order to maximize heat extraction and minimize induced seismicity. Our assumption is that each fracture process is characterized by spectro-temporal features and patterns that are not picked up by current signal processing methods used in seismology, but can be identified by the human auditory system and/or by machine learning. We present here a pluridisciplinary methodology aimed at addressing this problem, combining machine learning, auditory display and sound perception.

1. INTRODUCTION

Introduced in the 1960s for research purposes [1, 2], the transformation of seismic data into sounds has remained in use until the present day, though often for education or artistic use [3, 4, 5, 6]. However recent studies [7, 8, 9, 10, 11, 12, 13] have demonstrated the potential of auditory display in seismic research, using the power of the human auditory system to recognize patterns in signals, and to produce alternative signal descriptions that might provide seismologists with new insights and hypotheses. Concurrently, there is a recent trend in seismic research to use machine learning techniques [14, 15, 16, 17] for automatic classification purposes, bringing interesting results to the field of seismology.

Earthquakes are mainly caused by tectonic stress occurring in the Earth's crust and upper mantle. However, they can also be triggered by anthropogenic activity ("human-induced seismicity"). Human-induced earthquake may be the result of high pressure water injection into the ground, as it is found in Oklahoma during waste water disposal activity. Earthquakes can also be a by-product of human activities in geothermal fields, where water is allowed to percolate through the rock and pick up its heat, which is then used to generate electricity. Injecting under hydrostatic pressure cold water into a hot reservoir can drive different fracture mechanisms: hydraulic fracturing (microcracking driven by elevated fluid pressure), frictional sliding on existing faults (triggered by fluid pressure changes), and thermal cracking (driven

by changes in volume of the rock and fluid as their temperatures change). One current challenge in the field of seismology is to be able to identify these fracture mechanisms from the seismic signals, possibly in real time, so we can have control on them in order to maximize the heat extraction while minimizing the fluid pumping and induced seismicity.

In this project we focus on the earthquakes occurring at an active geothermal reservoir at "The Geysers" in Sonoma County, CA, USA¹.

2. METHOD

The main question is the identification of fracture processes from the physical measurement (*i.e.* seismic recordings). Our assumption is that each fracture process is characterized by spectro-temporal features and patterns that are not picked up by current signal processing methods used in seismology, but can be identified by the human auditory system and/or by machine learning. We are designing an experimental method for addressing this question, bringing together psychoacoustics and machine learning. To what extent can the human auditory system identify fracture processes? Can we teach computers to identify fracture processes?

2.1. Unsupervised machine learning

The first step of the method aims to reduce the large dataset to a smaller set of signals which are typical of this dataset (*i.e.* which sample effectively the dataset), so that they can be processed by human listeners in a reasonable amount of time. Indeed, earthquakes happen constantly at the Geysers, making the 3 year-long catalog of seismic data we focus on encompass more than 46,000 events. Machine learning techniques must be used in order to reduce the dataset to its most typical elements, in reasonable number, so that they can be given to listeners.

Because nothing is assumed about the patterns that are to be found in the signals, an unsupervised approach to machine learning is used. In particular the machine is not asked to extract particular features on which the clustering would be based (in the contrary to previous studies using machine learning technique in the field of seismology, *e.g.* [14]). Because the subset extracted via the machine learning is aimed at being transformed into sounds, the spectrogram representation of the signals is preferred to time series, assuming this representation of the data can be linked more easily to how humans perceive sound than raw waveforms. The unsupervised machine learning is implemented as follows:

¹<http://esd1.lbl.gov/research/projects/induced.seismicity/egs/geysers.html>



This work is licensed under Creative Commons Attribution Non Commercial 4.0 International License. The full terms of the License are available at <http://creativecommons.org/licenses/by-nc/4.0>

1. Learn patterns on the dataset through Non-Negative Matrix Factorization (NMF) [18, 19] on spectrograms. Each spectrogram can be reconstructed via the product of a matrix of activation coefficients (one for each spectrogram, smaller dimension than the spectrogram) and of a dictionary of patterns (a matrix common to all spectrograms);
2. Reduce the dimension of the activation matrices through Hidden Markov Models (HMM) [20];
3. Cluster the HMM-modeled activation matrices with the K-means method [20]. The machine is asked to define clusters of similar events (*i.e.* “close” vectors, according to a certain distance). Again, no hint is given to the machine on how to define these clusters. The number of clusters asked to the algorithm is defined by the user (see Sec. 3).

The only tuning that was done during the machine learning phase was the adjustment of the number of clusters. Having started arbitrarily with 10 clusters, we are investigating the effect of changing the number of clusters from 2 to 20.

The K-means algorithm assigns to each cluster a centroid. Spectrograms are then put in the cluster minimizing a distance to the centroid. This spectrogram/centroid distance can be used in order to assess the “typicality” of each spectrogram in a cluster: Spectrograms closer to the centroid are thought to be more typical of the cluster than further spectrograms. These prototypes will be selected for the listening tests, assuming that they are sufficient to provide a good and comprehensive overview of the clusters.

2.2. Audification of selected data

Seismic and sonic waves being very similar, the easiest and most conservative sonification technique is “audification” [21, 10, 11]. This technique consists in just changing the time-scale of the time series data (in our case doing a time-compression, equivalent to moving the infra-sonic frequency content of the seismic signal back up to the audible range, or to increasing the sampling frequency). Selected signals will be audified for the listening tests.

2.3. Listening tests

The interpretation of the clustering and of the criteria of “blind” machine categorization will be made through listening tests conducted with the audified data. Building on previous results that showed the potential of humans to assess audified seismic data [11, 13], the audified sounds will be the stimuli of a free sorting task. In such a task, the participants are asked to categorize sounds according to the perceived similarity. Each category then has to be verbally described, as in [13, 22, 23]. Such a categorization and verbalization task has a double purpose. First, we will check if machine and human categories agree, *i.e.* if items from the same machine-produced cluster are perceived as similar (grouped into the same category), and if items from different clusters are grouped into distinct categories. Second, the verbal description of each category gives us access to the criteria chosen by the listeners to proceed to the grouping of sounds. We will seek to interpret and link these criteria with audio descriptors (*e.g.* those defined in [24]) that can be computed on the spectrograms, but also on the spectra and waveforms, if applicable.

The abovementioned listening tests will be conducted both in a laboratory setting, and on a crowd-sourcing internet platform. Initial results from informal listenings demonstrate clear sonic differences among clusters.

2.4. Supervised machine learning

Once the signal features that are relevant to the listeners will be identified (and we assume that some of them will presumably correspond to features described in [13]), they will be computed for every signal in the dataset to check if categories produced by listeners can be retrieved by the machine. Eventually, these features — that are expected to be related to fracture processes or path effects — will feed a supervised machine learning onto another dataset: the algorithms will be fed with labelled data and will be trained to identify and categorize new data, based on criteria that are relevant from the rock mechanics point of view.

3. FIRST RESULTS

Our results from unsupervised learning show that if the number of clusters is higher than 4, the clustering comes up with one cluster gathering all events of higher magnitude. Such a high-magnitude cluster is not very informative, since the estimation of the magnitude of an event is a relatively easy task that does not need machine learning or auditory techniques. Therefore results are presented here with 2 clusters asked to the machine. These two clusters are referred to as C1 and C2 in the following.

The clusters are not characterized by any spatial criteria (they are not different faults), as Fig. 1 shows. In other words, seismic events are not clustered according to their location with respect to the source location (this would have been the case if *e.g.* one cluster had clustered all events north of the station). In a similar way, the clustering is not done according to the depth of the events (*i.e.* deep and shallow events are not separated out). However, the clustering turns out to be based on the occurrence time of the seismic events. Fig. 2 shows histograms of the occurrence times of events in each of the two clusters. This graphical representation shows distributions of dates that seem to slot very well together. This may suggest that similar physical phenomena are taking place during different periods of time. This may also be related to the history of water injection, potentially changing the fracture processes that are triggered, or the steam/liquid water balance in the rocks. We expect listening experiments to provide hints to help better understand the processes at work.

4. ACKNOWLEDGMENTS

Waveform data for this study were accessed through the Northern California Earthquake Data Center (NCEDC), doi:10.7932/NCEDC.

The authors thank Brian Bonner, Mark Warner, Kurt Nihei, Roland Gritto, Seth Saltiel, Heather Savage, Paul Johnson and Yen Joe Tan for fruitful discussion.

This project is funded through Columbia University RISE (Research Initiative in Science and Engineering) grant awarded to Ben Holtzman.

5. REFERENCES

- [1] S. S. Speeth, “Seismometer sounds,” *Journal of the Acoustical Society of America*, vol. 33(7), pp. 909–916, 1961.
- [2] G. E. Frantti and L. A. Levereault, “Auditory discrimination of seismic signals from earthquakes and explosions,”




Figure 1: Map of the events (dots) in the catalog, sorted by clusters as indicated by the color code. The algorithm was asked to produce 2 clusters, named C1 (events in C1 in black color) and C2 (events in C2 in black red). The size of the dots is proportional to the magnitude of the corresponding events. The seismic station is indicated with a green triangle in the middle of the area.

Bulletin of the Seismological Society of America, vol. 55(1), pp. 1–25, 1965.

[3] B. Holtzman, J. Candler, M. Turk, and D. Peter, Seismic Sound Lab: Sights, Sounds and Perception of the Earth as an Acoustic Space. London, UK: Springer Verlag, 2014, pp. 161–174.

[4] D. Kilb, Z. Peng, D. Simpson, A. Michael, M. Fisher, and D. Rohrlick, “Listen, watch, learn: Seissound video products,” Seismological research letters, vol. 83(2), pp. 281–286, 2012.

[5] M. Meier and A. Saranti, “Sonic explorations with earthquake data,” in Proc. of the 14th International Conference on Auditory Display (ICAD), Paris, France, 2008.

[6] Z. Peng, C. Aiken, D. Kilb, D. R. Shelly, and B. Enescu, “Listening to the 2011 magnitude 9.0 tohoku-oki, japan, earthquake,” Seismological research letters, vol. 83(2), pp. 287–293, 2012.

[7] C. Hayward, Listening to the Earth Sing. Boston, MA, USA: Addison-Wesley, 1994, pp. 369–404.

[8] F. Dombois, “Using audification in planetary seismology,” in Proc. of the International Conference on Auditory Display (ICAD), Espoo, Finland, 2001.

[9] —, “Auditory seismology: On free oscillations, focal mechanisms, explosions and synthetic seismograms,” in Proc. of the International Conference on Auditory Display (ICAD), Kyoto, Japan, 2002.

[10] L. Boschi, A. Paté, B. Holtzman, and J.-L. L. Carrou, “Can auditory display help us categorize seismic signals?” in Proc. of the 21st International Conference on Auditory Display (ICAD), Graz, Austria, 2015, pp. 306–307.

[11] A. Paté, L. Boschi, J.-L. L. Carrou, and B. Holtzman, “Categorization of seismic sources by auditory display: A blind test,” International Journal of Human-Computer Studies, vol. 85, no. 85, pp. 57–67, 2016.





Figure 2: Histogram of the time of occurrence (grouped by months) of events in clusters C1 (top) and C2 (bottom).

- [12] P. Dell'Aversana, G. Gabbriellini, and A. Amendola, "Sonification of geophysical data through timefrequency analysis: theory and applications," Geophysical Prospecting, pp. 1–12, 2016.
- [13] A. Paté, L. Boschi, D. Dubois, J.-L. L. Carrou, and B. Holtzman, "Auditory display of seismic data: On the use of experts' categorizations and verbal descriptions as heuristics for geoscience," Journal of the Acoustical Society of America, vol. Accepted on Feb. 15th, January 2017.
- [14] F. Provost, C. Hibert, and J.-P. Malet, "Automatic classification of endogenous landslide seismicity using the random forest supervised classifier," Geophysical Research Letters, vol. 44, pp. 113–120, 2017.
- [15] G. Curilem, J. Vergara, G. Fuentealba, G. A. na, and M. Chacón, "Classification of seismic signals at villarrica volcano (chile) using neural networks and genetic algorithms," J. Volcanol. Geotherm. Res., vol. 180(1), pp. 1–8, 2009.
- [16] F. Dammeyer, J. R. Moore, C. Hammer, F. Haslinger, and S. Loew, "Automatic detection of alpine rockslides in continuous seismic data using hidden markov models," J. Geophys. Res. Earth Surf., vol. 121, pp. 351–371, 2016.
- [17] P. B. Quang, P. Gaillard, Y. Cano, and M. Ulzibat, "Detection and classification of seismic events with progressive multi-channel correlation and hidden markov models," Comput. Geosci., vol. 83, pp. 110–119, 2015.
- [18] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," Nature, vol. 401, pp. 788–791, 1999.
- [19] —, "Algorithms for non-negative matrix factorization," in Proc. of the 13th International Conference on Neural Information Processing Systems, Denver, CO, USA, 2000, pp. 535–541.
- [20] C. M. Bishop, Pattern recognition and machine learning, ser. Information Science and Statistics. New York, NY, USA: Springer Verlag, 2006.
- [21] F. Dombois and G. Eckel, Audification. Berlin, Germany: Logos Verlag, 2011, ch. 12, pp. 301–324.
- [22] A. Paté, J.-L. Le Carrou, B. Navaret, D. Dubois, and B. Fabre, "Influence of the electric guitar's fingerboard wood on guitarists' perception," Acta acustica united with acustica, vol. 101, pp. 347–359, 2015.
- [23] C. Guastavino, "Categorisation of environmental sounds," Canadian journal of experimental psychology, vol. 61, pp. 54–63, 2007.
- [24] G. Peeters, B. L. Giordano, P. Susini, N. Misdariis, and S. McAdams, "The timbre toolbox: Extracting audio descriptors from musical signals," Journal of the Acoustical Society of America, vol. 130(5), p. 2011, 2902–2016.