

Short Note

The F -Detector Revisited: An Improved Strategy for Signal Detection at Seismic and Infrasonic Arrays

by Stephen J. Arrowsmith, Rod Whitaker, Charles Katz, and Chris Hayward

Abstract This short article explores and extends the adaptive detection algorithm recently developed by Arrowsmith, Whitaker, *et al.* (2008). In particular, this article highlights its application for seismic data, compares results for colocated seismic and infrasonic data, and assesses detector performance through comparison with analyst picks. We assess the adaptive detector by generating receiver-operating characteristic (ROC) curves, illustrating the trade-off between detection probability and false-alarm probability, and comparing the results with the conventional F -detector. The results show that the adaptive detector performs much better than the conventional detector for both seismic and infrasonic data by maintaining high detection probabilities while significantly decreasing false-alarm probabilities, illustrating that correlated noise is ubiquitous for both types of data. The effect of the adaptation window is illustrated and shown to be especially important for infrasonic data where diurnal variations in ambient noise levels are pronounced. A window choice of 1 hr (i.e., significantly less than 24 hr) is shown to be adequate for representing variations in ambient noise levels.

Introduction

Most people with some experience of data analysis can identify signals from noise remarkably easily. The problem of automating this process for handling large quantities of data is more complex and requires one to establish criteria (i.e., hypotheses) that can be tested numerically. Previous studies have typically implemented the simple hypothesis that noise is spatially uncorrelated between the individual elements in a seismic or infrasonic array. Signals are assumed to arrive as plane waves and can therefore be aligned via beamforming.

The difference between signal and noise depends on one's definition. In a practical monitoring perspective, the user is typically interested in relatively high-amplitude transient signals—not in low-amplitude continuous signals (which may be associated with a wide range of natural and anthropogenic sources). Under this scenario, the conventional null hypothesis (i.e., noise is completely uncorrelated) is inappropriate and commonly violated—leading to large numbers of detections from correlated noise (referred to here as clutter). Of course, if the source level of the transient of interest is inherently low, or if it is distant from the receiver array, then the transient event can become submerged into the low-level correlated clutter and be manifest simply as a brief excursion in the amplitude level of the overall ambient.

Historically, techniques employed to mitigate this problem have been somewhat *ad-hoc*, typically utilizing histor-

ical data at a given array to set detection thresholds that work well in practice. However, in addition to the obvious statistical limitations of this approach, such techniques cannot be easily exported to new arrays—especially arrays with no recording history. We therefore require a statistically robust algorithm that satisfies the following three criteria: (1) minimal need for manual tuning, (2) accounts for real ambient noise, and (3) ability to be applied operationally in near real time.

Arrowsmith, Whitaker, *et al.* (2008) outline a simple array-based detection algorithm that effectively satisfies all three criteria, which they then implement into a network-based infrasonic-monitoring algorithm. In this article, we apply the new detection algorithm to colocated seismic and infrasonic array data from Pinedale, Wyoming. We compare the results with conventional detectors by computing receiver-operating characteristic (ROC) curves, allowing us to directly evaluate detection and false-alarm probabilities for the different methods. Finally, we comment on differences in the detector implementation for seismic and infrasonic data.

Methodology

As outlined by Blandford (1974), the F -statistic is defined as:

$$F = \left(\frac{J-1}{J} \right) \frac{\sum_{n=n_0}^{n_0+(N-1)} [\sum_{j=1}^J x_j(n+l_j)]^2}{\sum_{n=n_0}^{n_0+(N-1)} (\sum_{j=1}^J \{x_j(n+l_j) - [\frac{1}{J} \sum_{m=1}^J x_m(n+l_m)]\}^2)}, \quad (1)$$

where J is the number of sensors, $x_j(n)$ is the waveform amplitude of the n th sample of the mean-free time series from sensor j , l_j is the time-alignment lag obtained from beamforming, n_0 is the start sample index for the processing interval, and N is the number of samples in the processing window. Shumway *et al.* (1999) showed that in the presence of correlated noise, the F -statistic is distributed as $cF_{2BT, 2BT(N-1)}$, where B is the bandwidth of filtered data in the processing window, T , N are the number of array elements, and c is given by:

$$c = \left(1 + N \frac{P_s}{P_n} \right), \quad (2)$$

with $\frac{P_s}{P_n}$ denoting the correlated-noise power to uncorrelated-noise power ratio (Shumway *et al.*, 1999). Ideally, the processing window, T , should be the same as or similar to the time duration of the transient signal, ΔT . In practice, this is not always the case, and if we would like to detect every signal present in the data we may need to process each record a few times using different processing time windows (and frequency bands). Arrowsmith, Whitaker, *et al.* (2008) showed how this theoretical framework could be implemented into a practical network-monitoring algorithm for infrasound data. In essence, we can estimate c by fitting the observed distribution of F -statistics in a given adaptive time window (w) to the theoretical distribution ($cF_{2BT, 2BT(N-1)}$) by finding the value of c that aligns the peaks of both distributions. We then apply a standard p -value threshold to find detections with a specified statistical significance. A new value of c is then computed for the next subsequent adaptive time window (w), allowing for adaptive compensation of temporally variable noise. We comment in detail on the choice of w in the following.

Constructing ROC curves is a standard practice for assessing detector performance (e.g., Johnson and Dudgeon, 1993). ROC curves allow us to evaluate the trade-off between the detection probability (P_D) and the false-alarm probability (P_F) for a range of detection thresholds. They also allow for direct comparison between the performances of different detectors, provided that all of the detectors are applied to the same set of input processes.

For real datasets, ROC curves can only be estimated because we must make *a priori* estimates of signal start and end times. In practice, this is very difficult and relies on the judgment of an analyst. However, as outlined in the Introduction, analyst picks are typically more reliable than automatic picks. Thus, it is reasonable to use analyst picks as *a priori* information to estimate ROC curves for automatic detection algorithms. For a given set of analyst picks, we can estimate P_D and P_F as follows:

$$P_D = \frac{\text{number of detected signals}}{\text{total number of signals}},$$

$$P_F = \frac{\text{number of noise detections}}{\text{total number of detection intervals during noise}}. \quad (3)$$

Test Dataset

In this study we have chosen to use data from colocated seismic and infrasound arrays in Pinedale, Wyoming (see the Data and Resources section and Fig. 1). This unique dataset allows us to directly compare detector performance for seismic and infrasonic data. We have selected one full day of data for this study: 6 November 2007 (Fig. 2). The choice of day is influenced by a previous study (Arrowsmith, Hedlin, *et al.*, 2008), which that found mining explosions from the Powder River basin could be detected infrasonically at Pinedale during the wintertime.

Results

Seismic and infrasonic data, shown in Figure 2, were picked by an analyst (in the 1–5 Hz frequency band). A total of 47 seismic events and 22 infrasound events were identified, for which start and end times were manually picked. The data were also processed with two detection algorithms: the previously described adaptive F -detector and the conventional F -detector (which assumes purely uncorrelated noise). In each case, we first apply a standard F - K analysis. An F - K spectrum is computed for a moving time window, providing the back azimuth and phase velocity of the maximum beam as a function of time. For each time window, the detection statistic (equation 1) is applied for the maximum beam (i.e., the values of l_j in equation 1 are directly obtained from the F - K analysis). Data were processed with the following parameters: time window is 10 sec, overlap is 50%, filter band is 1–5 Hz, and slowness parameters are -40 to 40 sec/deg (seismic) or -400 to 400 sec/deg (infrasound). For details on the implementation of a standard F - K analysis, the reader is referred to Rost and Thomas (2002). ROC curves for seismic and infrasound data were generated using equation (3) for the following p -value detection thresholds: 0.5, 0.6, 0.7, 0.8, 0.9, 0.91, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, and 0.99 (Fig. 3). We experimented with varying the adaptive time-window (w) results obtained for $w = 1$ hr and $w = 24$ hr that are shown in Figure 3.

An ideal detector would result in the following values: $P_D = 1$ and $P_F = 0$ (labeled in Fig. 3). The degree to which the ROC departs from the equality line ($P_D = P_F$), shown by the dashed lines in Figure 3, measures the distinctiveness

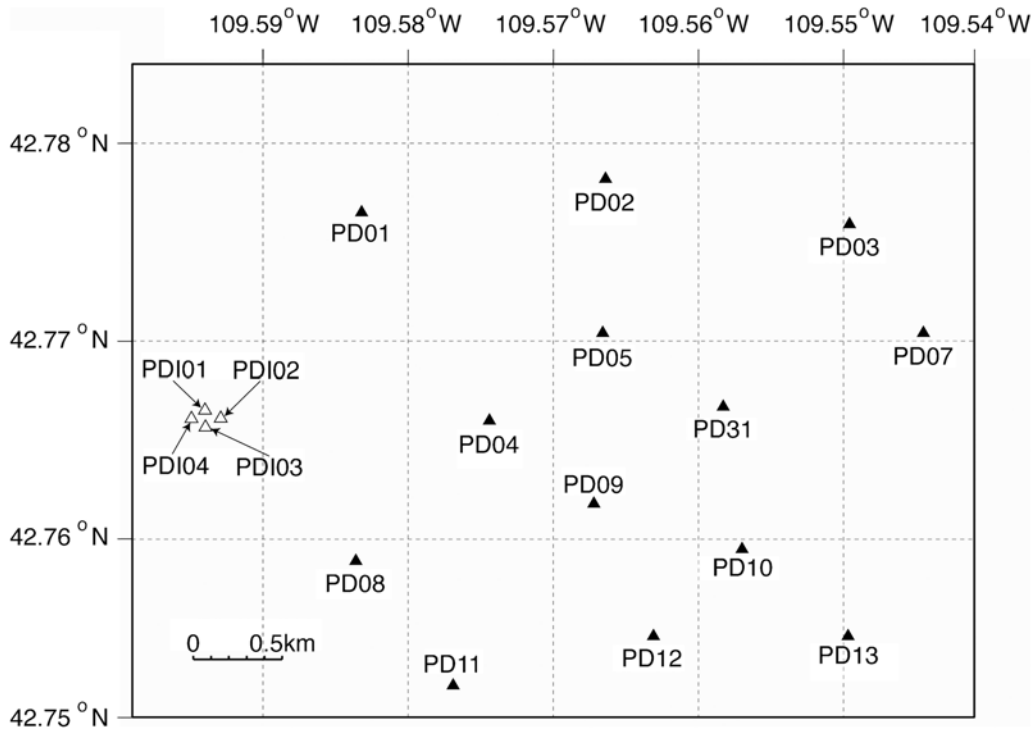


Figure 1. Map of the Pinedale seismic (filled triangles) and infrasound (open triangles) arrays.

of the two models ($H_0 = \text{noise}$ and $H_1 = \text{signal} + \text{noise}$). Taking the results for $w = 1$ hr first, it is clear that the adaptive detector performs much better than the conventional detector for both seismic and infrasound data. While both detectors have very high values for P_D , the conventional detector also has high values for P_F , indicating that large quantities of clutter are being detected. This is a direct consequence of the fact that the null hypothesis ($H_0 = \text{uncorrelated noise}$) for the conventional detector is being violated. In contrast, values of P_F for the adaptive detector are much lower (e.g., for seismic data, they are $P_F = 0.43$, $P_D = 1.0$ for $p = 0.5$ and $P_F = 0.02$, $P_D = 0.89$ for $p = 0.99$). Both seismic and infrasound results are very similar for the case where $w = 1$ hr. For a given p value, the corresponding values for P_D and P_F are comparable, demonstrating that the algorithm provides similar results for independent datasets at a given threshold and therefore requires minimal tuning.

The results obtained for $w = 24$ hr are very similar for the seismic data but quite different for the infrasound data. For the infrasound data, $P_F > 0.7$ for all adaptive detector results, indicating that the adaptive feature is not adequately compensating for the ambient correlated noise. The reason for this is clear from Figure 2: infrasonic noise levels vary significantly during the course of the day, with particularly high levels of incoherent noise during early-mid-afternoon (local time) due to solar-driven convection. Figure 4 illustrates the effect of this variation in noise on the calculated F -distributions in two time intervals—04:00–08:00 coordinated universal time (UTC) (late-evening local time) and

20:00–00:00 UTC (mid-afternoon local time). For the seismic data, the F -distributions in both time intervals are very similar. In contrast, for the infrasound data, the F -statistics are significantly higher in the evening than in the afternoon. Strong, incoherent solar-driven turbulence swamps coherent noise on the infrasound array, lowering the ambient F -statistics.

Clearly, for infrasound data, it is important that w be set significantly less than 24 hr in order to account for diurnal variations in the ambient noise field. In contrast, for seismic data this is less important, although it must be emphasized that seismic stations at Pinedale Array are located in boreholes and are therefore less sensitive to diurnal variations in the atmosphere. Of course, both type of sensors are subject to regular time-dependent ambient variations due to cultural activity in instances where the sensor arrays are proximate to human enterprise, for example, highways, factories, development infrastructure such as pumping stations and powerplants, etc.

Discussion and Conclusions

The foremost criterion for any automatic detector is to assess how successfully it can match human analyst picks. Our goal is to develop a detector for finding transient, relatively high-amplitude signals, which perform similarly to a human analyst (effectively removing the need for a human analyst in processing large quantities of data). As shown in this article, conventional detectors, such as the standard F -detector, do not provide this capability. The analysis performed in this article provides a direct comparison between




Figure 2. Seismic (top) and infrasonic (bottom) array data recorded at Pinedale on 6 November 2007. Each waveform is band-pass filtered from 1 to 5 Hz. Amplitudes are normalized by the maximum amplitude on each trace. Boxes enclosed by dashed lines highlight the time windows used for computing F distributions shown in Figure 4.

human analyst picks and automatic picks for the adaptive F -detector and for the conventional F -detector. Because this analysis assumes that we know signal start and end times *a priori*, it should be emphasized that, in this study, we are only estimating the ROC curves. However, this analysis confirms that the adaptive F -detector performs very similarly to a human analyst, whereas the conventional F -detector performs comparatively poorly for both seismic and infrasound data.

In the Introduction, we identified three additional criteria of importance in developing our automated detection algorithm. The first criterion, that our detection algorithm require minimal tuning, is satisfied in two ways: (1) the results shown in Figure 3 demonstrate that for a typical range of p -value thresholds used in hypothesis testing (i.e., 0.01–0.05), the values of P_D and P_F are comparable ($P_D > 0.87$, $P_F < 0.15$) for independent seismic and infrasound datasets (assuming w is set appropriately); and (2) the detector adapts to variable ambient noise, removing the need to manually

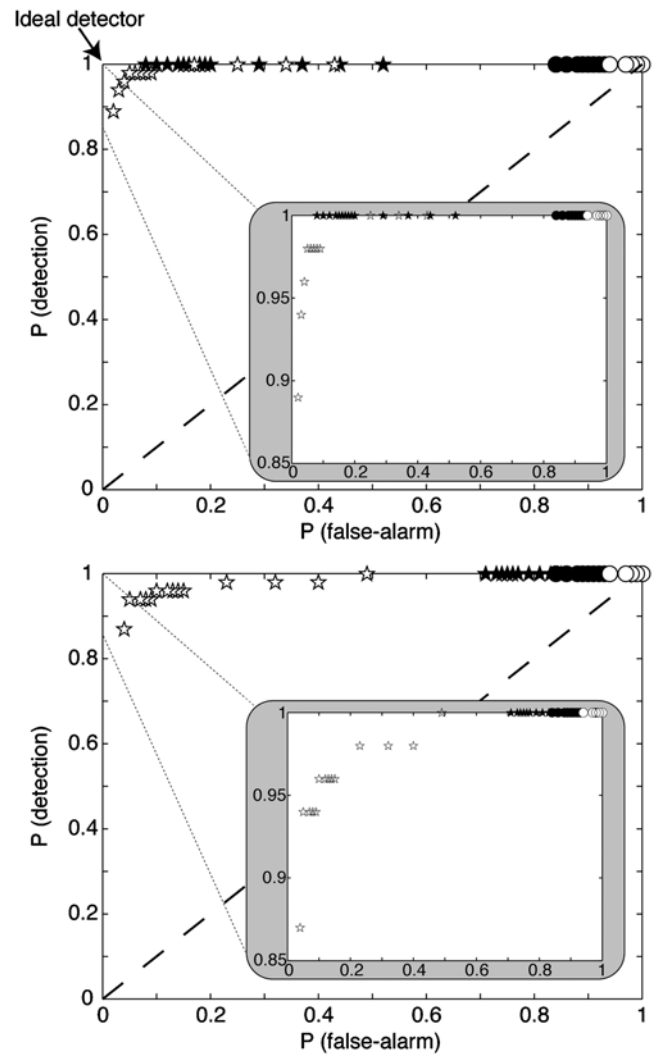


Figure 3. ROC curves for adaptive (stars) and conventional (circles) detectors for both infrasound (filled symbols) and seismic (open symbols) data. The top panel is for a long-term window (w) of 1 hr and the bottom panel is for $w = 24$ hr. The equality line ($P_D = P_F$) is shown by a dashed line on each panel. The inset plots show zooms on P (detection) from 0.85 to 1.

adjust thresholds in order to account for changes in noise. The second criterion, that the algorithm can adapt to temporally variable noise, is shown to be satisfied if the adaptive window length is set short enough for infrasound data in order to avoid diurnal variations in noise levels. Finally, although not specifically mentioned in the text, the third criterion (that the algorithm can be applied operationally in near real time) is satisfied. The adaptive F -detector requires minimal additional processing in comparison with the conventional F -detector.

Data and Resources

Seismic data used in this study can be obtained from the IRIS Data Management Center at www.iris.edu (network code: IM). Infrasound data used in this study were obtained

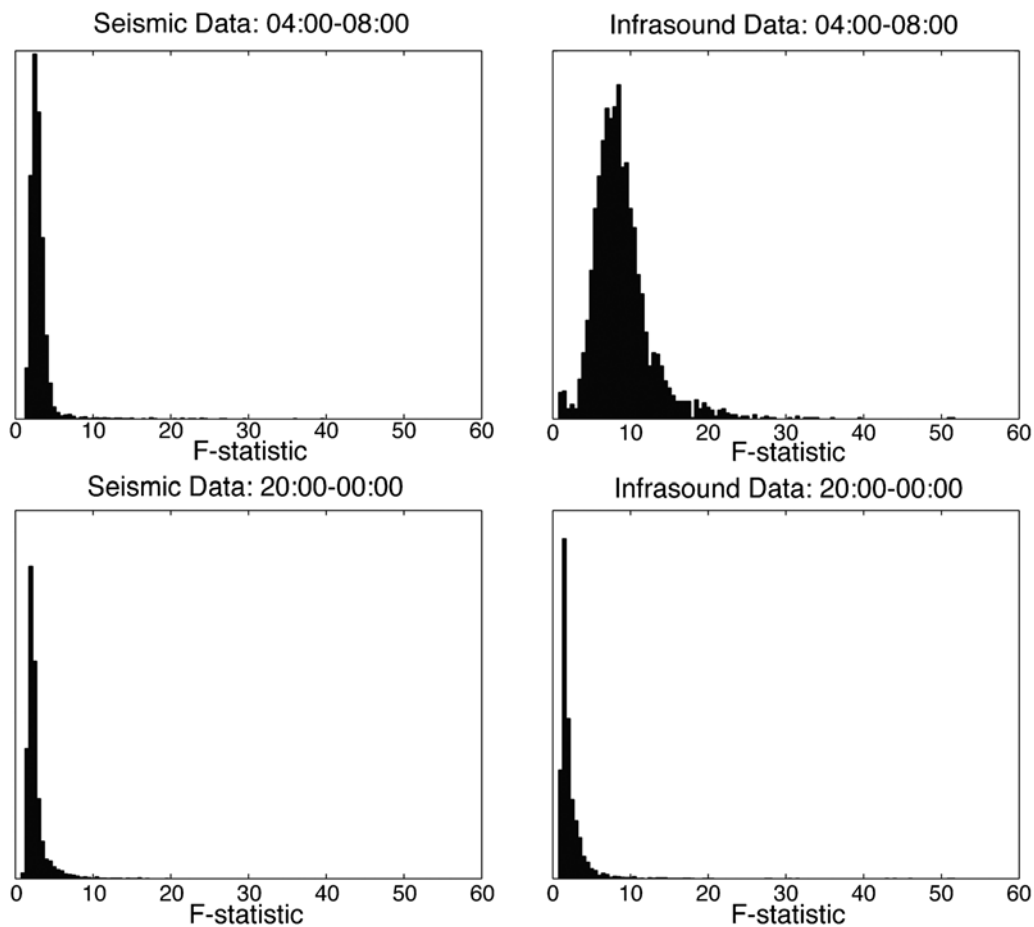


Figure 4. Distributions of raw F -statistics (from equation 1) for seismic (left column) and infrasonic (right column) data in two coordinated universal time intervals: 04:00–08:00 (top row) and 20:00–00:00 (bottom row).

as part of a temporary deployment run by Southern Methodist University and are not freely available to the community.

Acknowledgments

We thank George Randall, Steve Taylor, Diane Doser, and an anonymous referee for their comments on an earlier draft of this article. This work was supported in part by the U.S. Department of Energy, Office of Non-Proliferation Research and Development.

References

- Arrowsmith, S. J., R. Whitaker, S. R. Taylor, R. Burlacu, B. Stump, M. Hedlin, G. Randall, C. Hayward, and D. ReVelle (2008). Regional monitoring of infrasound events using multiple arrays: application to Utah and Washington State, *Geophys. J. Int.* **175**, 291–300, doi: 10.1111/j.1365-246X.2008.03912.x.
- Arrowsmith, S. J., M. A. H. Hedlin, B. W. Stump, and M. D. Arrowsmith (2008). Infrasonic signals from large mining explosions, *Bull. Seismol. Soc. Am.* **98**, 768–777.
- Blandford, R. R. (1974). An automatic event detector at the Tonto Forest Seismic Observatory, *Geophysics* **39**, 633–643.
- Johnson, D. H., and D. E. Dudgeon (1993). *Array Signal Processing: Concepts and Techniques*, Prentice-Hall, Englewood Cliffs.
- Rost, S., and C. Thomas (2002). Array seismology: methods and applications, *Rev. Geophys.* **40**, no. 3, 1008, doi 10.1029/2000RG000100.
- Shumway, R. H., S.-E. Kim, and R. R. Blandford (1999). Nonlinear estimation for time series observed on arrays, in *Asymptotics, Nonparametrics and Time Series*, S. Ghosh (Editor), Marcel Dekker, New York.

EES-17
Los Alamos National Laboratory
P.O. Box 1663
Los Alamos, New Mexico 87545
(S.J.A., R.W.)

Science Applications International Corporation (SAIC)
10260 Campus Point Drive
San Diego, California 92121-1578
(C.K.)

Department of Geological Sciences
Dedman College
Southern Methodist University
P.O. Box 750395
Dallas, Texas 75275
(C.H.)