The Optimal Correlation Detector?

Steven J. Gibbons[1]

[1] *Norges Geotekniske Institutt, Sognsveien 72, 0855 Oslo, Norway,*

E-mail: steven.gibbons@ngi.no

SUMMARY

Correlation detectors are now used routinely in seismology to detect occurrences of signals bearing close resemblance to a reference waveform. They facilitate the detection of low-amplitude signals in significant background noise that may elude detection using energy detectors, and they associate a detected signal with a source location. Many seismologists use the fully normalized correlation coefficient C between the template and incoming data to determine a detection. This is in contrast to other fields with a longer tradition for matched filter detection where the theoretically optimal statistic C^2 is typical. We perform a systematic comparison between the detection statistics C and $C|C|$, the latter having the same dynamic range as C^2 but differentiating between correlation and anti-correlation. Using a database of short waveform segments, each containing the signal on a 3-component seismometer from one of 51 closely spaced explosions, we attempt to detect P- and S- phase arrivals for all events using short waveform templates from each explosion as reference signals. We present empirical statistics of both C and $C|C|$ traces and demonstrate that $C|C|$ detects confidently a higher proportion of the signals than C without evidently increasing the likelihood of triggering erroneously. We recall from elementary statistics that C^2, also called the coefficient of determination, represents the fraction of the variance of one variable which can be explained by another variable. This means that the fraction of a segment of our incoming data that could be explained by our signal template decreases almost linearly with $C|C|$ but diminishes more rapidly as

C decreases. In most situations, replacing C with $C|C|$ in operational correlation detectors may improve the detection sensitivity without hurting the performance-gain obtained through network stacking. It may also allow a better comparison between single-template correlation detectors and higher order multiple-template subspace detectors which, by definition, already apply an optimal detection statistic.

Key words: Time series analysis, Earthquake monitoring and test-ban treaty verification, Induced seismicity

1   INTRODUCTION

Realizing that closely spaced seismic events generate similar signals (e.g. Geller & Mueller 1980) motivated the use of template-matching for a combined detection and identification of seismic signals. The fully normalized correlation coefficient, C, between two time-series had been used as a measure of waveform similarity, both for measuring accurately the time-delays between two signals (e.g. Van-Decar & Crosson 1990; Shearer 1997) and for providing a heuristic alias for inter-event distances (e.g. Menke 1999; Pirli et al. 2011). Traces of C, evaluated between a fixed waveform template and consecutive windows of an incoming data-stream, are functions of time which reach maximum values at samples when the segment of the incoming data best resembles the applied template, with lower values at other times. If the contrast between the maximum value and the background level of C is sufficiently high for a given sample, we can hypothesize that a new occurrence of a signal resembling our template has been found: a signal-specific detector. Early demonstrations of seismic correlation detectors on multi-channel data (Gibbons & Ringdal 2006; Shelly et al. 2007; Schaff 2008) pointed out a useful property of C; since the signals from two co-located seismic events are separated in time by an identical delay for all stations, a zero-moveout stacking of C traces will always add constructively at the time of matching and interfere both constructively and destructively at other times. This results in an almost perfect gain in signal-to-noise ratio (SNR), since the (correlation) signal is perfectly coherent and the (correlation) noise largely incoherent. The very high performance gain using seismic arrays and networks partly explains the increase in the application of correlation detectors and associated algorithms in recent years, as large and dense monitoring networks have become common for monitoring microseismicity (e.g. Benz et al. 2015; Duboeuf et al. 2021).

Seismology was relatively late in adopting template matching for the purposes of signal detection. This can largely be explained by the relatively limited circumstances under which we can know the temporal shape of a seismic signal a priori. The issues of applicability are discussed by Harris

et al. (2012) who review progress in both empirical and model-based signal processing, expressing doubt that tomographic models alone will extend template-matching with synthetic signals into the frequency band of relevance for detecting small earthquakes and explosions. In regions of clustered seismicity, where historical signals are often available, the picture is more encouraging (e.g. Richards et al. 2006) although Slinkard et al. (2013) demonstrate that even single aftershock sequences from large earthquakes are often poorly represented by small numbers of waveform templates. The field of image processing uses template-matching extensively and in multiple directions (e.g. Brunelli 2009) and the seismologist can exploit optimized image processing routines to find occurrences of a repeating signal in a datastream using only a few lines of python code with the packages `obspy` (Beyreuther et al. 2010), `numpy` (Harris et al. 2020), and `scikit-image` (van der Walt et al. 2014).

The shape of a seismic signal is determined by Earth structure, typically at spatial scales that are difficult to model. In other branches of physics, the shapes of signals are often more easily determined by theory and the matched filter (Turin 1960; van Trees 1968) has extensive applications in identifying signals in potentially noisy data. Acoustical signals underwater (Parvulescu 1995) and (at local distances) in the atmosphere (Taylor et al. 2013) have been shown to be amenable to detection using matched filters, as have, among many others, neurological time-series (e.g. Stamoulis & Richardson 2010) and gravitational waves (e.g. Abbott et al. 2016). Also in seismology, empirical signal detection algorithms with a robust basis in detection theory have been developed (Harris 2006; Harris & Paik 2006; Harris & Dodge 2011). The so-called subspace detectors are multiple dimensional empirical signal detectors which can allow a degree of variability in the detected signal. As Harris (2006) demonstrates, the correlator is the 1-dimensional subspace detector: the most sensitive, but also the most signal-specific. Increasing the dimension of the subspace detector broadens the geographical footprint over which the detected events can be located at the expense of a reduction in sensitivity and an increased probability of triggering on noise or unrelated signals. An extensive set of studies have now applied subspace detectors to diverse systems of clustered seismicity (e.g. Junek et al. 2015; Barrett & Beroza 2014; De La Hoz et al. 2021). In the subspace detector formalism, the detection statistic employed by the 1-dimensional correlator is equivalent to C^2.

`EQcorrscan` (Chamberlain et al. 2018) is a Python-based software system, free to use and easy to implement, which is rapidly becoming a workhorse for efficient and sensitive empirical signal detection in seismology (e.g. Warren-Smith et al. 2018; Hicks et al. 2019; Chamberlain et al. 2020; van Wijk et al. 2021). EQcorrscan currently employs the Gibbons & Ringdal (2006) prescription of calculating C traces for individual channels, prior to a stacking operation, although a subspace-detector implementation is also under development. Fast Matched Filter or FMF (Beaucé et al. 2018) is another system for large-scale correlation detector calculations, which also employs C as a basis

for detection. Gibbons & Ringdal (2012) presented a multi-channel correlation detector but with a slightly modified detection trace: $C|C|$. The function $C|C|$ has the same dynamic range as C^2 but varies in sign like C, such that the alignment in time of the detector traces can be analyzed using the post-processing method of Gibbons & Ringdal (2006). This proved to be crucial for the elimination of false alarms. Since these detectors were being run primarily on small to medium aperture seismic arrays, it was quite common to trigger on incoming wavefronts from different directions when the detection threshold was set very low. Such false alarms could easily be identified by array processing of the correlation coefficient traces. The $C|C|$ trace was later used by Ford & Walter (2015) to monitor the North Korean nuclear test site down to very low seismic magnitude.

Gibbons et al. (2014) demonstrated the extent of the moveout of the local maxima of the correlation coefficient traces to stations in different directions which results from differences in the relative locations of events in the source region. This needs to be corrected for when stacking over a network, in which the stations cover a wide range of azimuths from the source, and can be exploited to combine the detection and location processes. This is exactly what is done in the Match and Locate procedure (Zhang & Wen 2015). The study of Warren-Smith et al. (2018) also facilitates a degree of flexibility in the source location by allowing a degree of moveout between peaks in the correlation functions. Another recent and more flexible strategy is the Multisegment Cross-Correlation Approach of Gao & Kao (2020) (applied by Hughes et al. 2021, in generating an extensive seismicity catalog). These procedures increase greatly the applicability of matched signal detectors and enhance the information that can be gleaned from a detection. Another issue in the specification of empirical signal detectors is the question of individual-channel or multiplex data? Maintaining individual-channel detection statistic traces, and allowing moveout, likely provides greater flexibility at the possible expense of decreased efficiency. The multi-station validation method (Slinkard et al. 2014) could be used to maintain a low false-alarm rate on multiple detectors, each using multiplexed data.

In a classical detection process, the SNR is quite straightforwardly related to the energy in the signal and the energy in the background noise. The detection statistic in a matched filter detector relates only to the degree to which the detected signal matches the waveform template. A high-SNR signal will result in a diminished matched filter detection statistic if the detected signal is poorly represented by the template. This can even be the case for co-located events with a large discrepancy in the magnitude (e.g. Gibbons et al. 2018; Bachura & Fischer 2019) with larger wavelengths associated with the larger event. Carmichael & Hartse (2016) quantify the degradation of detector performance with partial signal decorrelation and Carmichael (2016) introduces the so-called Cone Detector to accommodate a deterministic departure from the shape of the template waveform. It is pointed out that the detection performance is diminished relative to that of the correlation detector.

In this study we limit ourselves to a single-template multi-channel correlation detector and address one specific question; is there a consequential difference in performance between the traces C, C^2, and $C|C|$? All these traces are used in current operational matched signal detectors in seismology and it would be beneficial to understand how using one or the other affects detection performance. We exploit the dataset of signals from almost co-located explosions provided by Gibbons et al. (2020) as (a) they are readily available and fully documented, (b) the dataset contains signals from known explosive sources with small but measurable, and easily understood, waveform differences from event to event, and (c) there is exactly one signal in each short data segment. In section 2 we review exactly what we are referring to when we discuss correlation detectors and demonstrate the detectability of signals with detection statistics C and $C|C|$. In section 3 we perform a systematic comparison between these detection statistics over the chosen dataset and, in section 4, we seek to understand any difference in performance in terms of elementary results in statistics. Finally, in Sec. 5 we draw conclusions and make recommendations regarding implementation and subsequent evaluation.

2 DEFINITIONS AND FORMULATION

The detector of Gibbons & Ringdal (2006) calculated for each channel i in a seismic array or network a fully normalized correlation coefficient trace

$$C_i(t) = \frac{\boldsymbol{x}_i(t_0).\boldsymbol{y}_i(t)}{\sqrt{\boldsymbol{x}_i(t_0).\boldsymbol{x}_i(t_0)}\sqrt{\boldsymbol{y}_i(t).\boldsymbol{y}_i(t)}} \tag{1}$$

where $\boldsymbol{y}_i(t)$ is a vector of M time samples of the incoming data stream starting at time t and $\boldsymbol{x}_i(t_0)$ is the fixed template vector of length M samples starting at time t_0. We can save ourselves the computation of the $\sqrt{\boldsymbol{x}_i.\boldsymbol{x}_i}$ term at each time, t, by normalizing the template vector such that

$$\boldsymbol{x}_i(t_0).\boldsymbol{x}_i(t_0) = 1 \tag{2}$$

a priori.

Fig. 1 displays the concept of detecting a signal from one seismic event using a waveform template extracted from a different seismic event on a single 3-component station. Both signals were generated by surface explosions carried out by the Finnish military at a distance of approximately 208 km from the station. The open access study of Gibbons et al. (2020) provides both the Ground Truth coordinates of the explosions and the waveform segments (in miniseed format) used to calculate the traces displayed. The event (H01) from which the template is obtained is about 275 meters closer to the station than the detected event (H02). A fully normalized correlation trace $C_i(t)$ is calculated for the three components KEV.BHZ, KEV.BHN, and KEV.BHE and the correlation trace displayed, $C(t)$, is the mean of these with

$$C(t) = \frac{1}{n_c} \sum_{1}^{n_c} C_i(t) \qquad\qquad (3)$$

where n_c is the number of channels: in this case, 3.

Several correlation traces $C(t)$ are displayed in Fig. 1, each one calculated using a waveform template, starting at the time indicated, with a duration in seconds as indicated. As we are only using the short data segments provided in the supplementary material to Gibbons et al. (2020), these correlation coefficient traces become shorter as the length of the template increases. If there are N samples of the incoming data, and M samples in the template, then the correlation traces have length $(N - M + 1)$ samples. In the trace at the top, the template is only one second long: a vector of 40 samples, on each of the three components. The maximum value of the detection trace $C(t)$ is 0.928, although the variability of $C(t)$ is large. As all the waveforms are bandpass filtered 2-8 Hz, the one-second long template contains only a few cycles for each channel and the likelihood of near-coincidental phase-alignments with subsequent 1-second lengths of data is relatively high. The maximum value is attained at the time of the P-wave arrival but a secondary peak a few seconds later places a clear constraint on where the detection threshold would have to be set to avoid multiple detections. Had this been an earthquake sequence, we may not have been able to rule out an aftershock a few seconds later. However, these signals are generated by well-constrained explosions and we can rule out secondary events.

The second trace from the top uses a two-second long waveform template on each channel and has twice as many samples which need to be matched. The maximum value decreases slightly (to 0.892) but the background level decreases more. The secondary peak in the top trace is not significant in the second trace. In the third trace, the length of the template is doubled again. The maximum correlation value decreases marginally to 0.830 and the background level of the detection traces is reduced further. It is worth looking more closely at the template T04. The direct P-wave (in the first two seconds) has a significantly higher amplitude than the later two seconds of P-wave coda for the vertical channel. While there are twice as many samples in the T04 template than in the T02 template, it is clear from Eq. 1 that, for the vertical component at least, the new samples contribute less to the similarity measure than the higher-amplitude direct P-phase samples. Since the explosions are almost due south of station KEV, the BHN (horizontal) channel is almost identical to the radial rotation of the horizontal channels. The shape of the radial component is typically very similar to the shape of the vertical channel close to the P-arrival and this is reflected in the detection traces on the BHZ and BHN channels. The BHE (horizontal) channel is almost the transverse rotation for this particular source and station combination, and the P-arrival is as expected far less impulsive; the P-arrival and coda have similar amplitudes on this trace. We note that the ratio-to-moving average approach of Gibbons et al.

(2012) and the multiple segment approach of Gao & Kao (2020) are both designed to mitigate the effect of amplitude variability in the wavetrain.

In trace 4 (T08) the length of the template is again doubled and a far longer segment of coda is considered. Gibbons et al. (2020) used templates of length 6s for all arrivals since the S-arrival at the closest station in that study arrived some 7s after the P-arrival. Six seconds was therefore the longest template length that could be used on all stations for both P- and S- arrivals without combining P- and S-waves in a single template. Both templates T08 and T15 extend well into the P-coda at KEV, without reaching the S-wave. While not necessarily a significant issue for signal detection, the length of the template taken for measuring a time-delay may influence the interpretation of a measurement. Using a longer template will give a clearer peak but the user needs to bear in mind that slower propagating coda waves have contributed to this measurement. The variance of the time-delay estimate may have been reduced but at the cost of a greater bias. The final two templates in Fig. 1, T30 and T60, both include the more slowly propagating S-wave arrival and varying lengths of S-wave coda. If the two events are not exactly co-located then the contributions to the detection trace from the P-wave and S-wave segments will compete, potentially constructively and potentially destructively, depending both upon the hypocentral locations relative to the stations and the properties of the wavetrains. The reduction of the maximum value of C continues as the length of the template increases but the suppression of the background level of C becomes even more visibly effective. The matched filter detector now has a large time-bandwidth product and the requirements that must be met to generate a statistical outlier in the traces T30 and T60 are significantly higher than in the T01 or T02 traces.

The detector described in Gibbons & Ringdal (2006) was an ad hoc procedure, converted from an earlier code which measured time-differences between similar signals, adapted specifically to detect signals from small explosions using time-domain cross-correlation (Stevens et al. 2006). A more efficient and purpose-designed code was later developed and applied by Gibbons & Ringdal (2012). This code performed correlations far more rapidly in the frequency domain and used a slightly modified detection statistic

$$\frac{(\boldsymbol{x}_i(t_0).\boldsymbol{y}_i(t))abs(\boldsymbol{x}_i(t_0).\boldsymbol{y}_i(t))}{\boldsymbol{y}_i(t).\boldsymbol{y}_i(t)} = C_i(t)|C_i(t)| \tag{4}$$

where we simply demand beforehand that the template vector, $\boldsymbol{x}_i(t_0)$, is unit-normalized. As stated by Gibbons & Ringdal (2012), the change to this detection statistic was motivated by efficiency; without the need to compute the square root terms in the denominator, this term consists only as a product and a quotient of two convolutions. The behaviour of the detection statistic $C|C|$ was not compared explicitly with the behaviour of C.

Fig. 2 displays the correlation traces associated with an attempt to detect the S-phase for event H04 in Gibbons et al. (2020) using a 2s long template taken from the S-phase from event H03. With

40 Hz sampling, the template consists of 80 samples and is significantly shorter than the full wavetrain (which is slightly in excess of a minute). We display both single-channel C and $C|C|$ traces and the corresponding 3-component stacks. It is typically the stacks that are used for detection purposes, but detailed visualization of the single-channel traces can provide insight into the behaviour of the stack traces. The detection of S-phases is typically more complicated than the detection of P-arrivals since the signal arrival is competing not just against the ambient background noise but also the possibly very energetic P-wave coda. The P-wave coda does not cease when the S-wave arrives and we typically need to find a filtering of the wavefield which maximizes the contrast between the waveform at the S-wave arrival and the underlying coda-waves. The dissimilarity of P-waves from two closely spaced explosions is almost entirely due to the small differences in the path taken. The S-waves in addition will have slightly different time-offsets relative to the P-wave coda which will likely increase the dissimilarity.

The uppermost channels in Fig. 2 panels b) and c) show the 3-component stacks for the C and $C|C|$ traces respectively. The ratio of the maximum value to the background level is clearly higher for the $C|C|$ stack in panel (c) than for the C stack in panel (b). For the single-channel correlation traces, the picture looks somewhat more complicated with what appears to be a greater number of outliers for the single-channel $C|C|$ traces, although the averaging over the three components appears to mitigate the effect of this variability. It is to be noted that the second most significant peak in panel (c) for the S-wave template comes for the P-wave arrival. This secondary peak is most significant on the BHE component; it would be a clear detection on this single channel. However, there is no corresponding peak on the BHN and BHZ channels and the peak is diminished in the average stack (labelled AVGC2). The secondary peak at the time of the P-arrival is a false detection in the sense that it is a different part of the wavetrain that exhibits spurious similarity to the template. Just as the main (S-wave) $C|C|$ peak exceeds the $C|C|$ background level (panel c) more than the corresponding C peak exceeds the C background level (panel b), the spurious P-wave $C|C|$ detection is rather more visible than the corresponding C detection. Whether or not such effects would lead to a different ratio of true positives to false positives is not clear without rigorous experimentation.

The detection is visually more convincing for the traces in panel (c) than for those in panel (b), but we need to evaluate the distribution of the detection statistics over a much broader set of calculations than those displayed in Fig. 2 to be able to draw safe conclusions. Since you can convert between individual values of C and $C|C|$ using a simple transformation, it would be tempting to assume that a threshold value τ for C would correspond to a threshold value τ^2 for $C|C|$. However, one issue is that we are finding detections on a stack (in this case for 3 components but, in principle, also possibly on an arbitrarily large number of sensors) and it is not clear how the cancellation of positive and

negative values differs between the two different detection traces. Another issue is that a fixed value threshold is rarely used; we almost always apply a threshold which is a function of some statistic of the background value. The effectiveness of a detection statistic will depend both on the value at the time of a true signal match and the chosen metric of the background level.

3 A SYSTEMATIC COMPARISON OF TWO DETECTION TRACES

In the study of Gibbons et al. (2020) there are 51 explosions for which we have a signal on the 3-component station KEV. This means that, if we take a waveform template from one event, there are 50 other explosions for which we can attempt to detect the signal using our template. For a single portion of the KEV wavetrain, there are then 2550 correlations we can perform using only the short 3-component waveform segments contained in this dataset. This is to say that we never attempt to detect the same signal from which the template was extracted. However, if we attempt to detect a signal from event b using a template from event a then we also attempt to detect a signal from event a using a template from event b. We attempt to detect P-waves and S-waves for each explosion in separate calculations, with waveform templates of length 2s starting at the times specified in the supplementary material of Gibbons et al. (2020).

We have a total of 102 3-component templates of length 2s, one for the P-arrival and one for the S-arrival from each of the 51 explosions. For each such template, we attempt to detect the corresponding section of the wavetrain from each one of the other explosions using a 3-component stack matched signal detector on the full 5-minute waveform segments. For each template and "incoming data" waveform segment, we calculate two separate sets of matched signal detection traces: C and $C|C|$, as displayed in Fig. 2. We calculate the single-channel detection traces and then form the 3-component stack. For each 3-component stack, we identify the maximum value of the detection statistic trace which falls within 2.5s of the anticipated time of the sought arrival (using the times provided by Gibbons et al. 2020). We denote this value the "Target Maximum Value". We then find the maximum value of trace excluding the target window extending 2.5s on either side of the anticipated arrival time. We will refer to this value as the "Highest Noise Value". While the standard deviation (or any other statistic of the background level of the detection trace) could provide the basis for an SNR for the detection statistic, it is this highest noise value which would represent a false alarm: a trigger on a waveform segment which does not contain a target waveform. All values in the resulting 3-component stack are then reordered from lowest to highest and we record the values at the percentiles 1, 5, 10, 25, 50, 75, 90, 95, and 99%.

A selection of these distributions are displayed in Fig. 3. We select the first 350 calculations out of the 2550 possible for display purposes only. Every line in each of the panels in Fig. 3 indicates

the distribution of values for this 5-minute long detection trace according to the key below. Each such line is a visual summary of the 3-component stack traces as shown in Fig. 2. Panels (a) and (b) of Fig. 3 display the distributions of values for the 3-component stacks when using P-wave templates and the C and $C|C|$ traces respectively. Panels (c) and (d) show the equivalent distributions when S-wave templates are used.

The target maximum values for P-wave templates (panels a and b) exceed the labelled percentiles of the background levels and the "highest noise values" for all calculations. For the S-wave templates, the target maximum values exceed the same thresholds for what appears to be the majority of cases. However, there are a number of calculations for which the target maximum value is very close to, or even lower than, the "highest noise value" such that a confident detection would not be possible without also admitting one or more false detections. We note that the distributions for the background levels and "highest noise values" are indistinguishable between the P-wave and S-wave templates. The templates all have the same length, and are bandpass-filtered between the same frequencies, and the likelihood of coincidentally matching an S-wave template appears no different in this case to the likelihood of coincidentally matching a P-wave template. The target maximum values of the distributions are however typically lower for the S-wave templates than for the P-wave templates. This follows from the higher variability in the shape of the S-waves than in the shape of the P-waves. For some S-waves, with the short templates applied, the target maximum value is lower than the highest noise value; in practice, this means a non-detection.

In the $C|C|$ traces, all values are drawn closer to zero than in the C traces; this is a direct consequence of the fact that all values of C lie between -1 and 1. What is crucial for any difference in performance of a detector is a change in the relationship between the "positive match" values and the diverse metrics of the background levels. Fig. 4 displays ratios between the target maximum (or "positive match") values and the indicated metrics of the background levels. In Fig. 3 we saw a scatter of the detection values themselves for a selected subset of the calculations. In Fig. 4 we see the critical ratios only for all of the calculations. To make the figure easier to read, we rearrange the values in each curve from smallest to largest. In panel (a) we consider the ratio between the target maximum value and the standard deviation of the background values. It is clear that we would need to set a threshold of many times σ before we failed to trigger on P-wave, with both the C and $C|C|$ detection traces. Even though the values of $C|C|$ are lower than the corresponding values of C, the standard deviation of the background level diminishes even faster resulting in a higher ratio. The more significant picture is displayed in panel (b): the ratios between target maximum values and the highest noise values (i.e. peaks that do not correspond to the times of the target signal). An arbitrary transformation could be applied to the correlation coefficient traces to stretch out the distance between relatively high values

and relatively low values. Such a transformation would of course be of no advantage if it resulted in a significantly higher number of false alarms.

Only 1 out of 2550 calculations (0.039%) for detecting P-waves using the $C|C|$ trace has a ratio of target maximum value to highest noise value below 2 (ratio = 1.99). 361 out of 2550 calculations (14.16%) for detecting P-waves using the C trace have ratios of target maximum value to highest noise value below 2 (lowest ratio = 1.41). For the calculations to detect S-waves using the $C|C|$ trace, 51 of 2550 calculations (2.0%) have target maximum to highest noise value ratio below 1 (i.e. clear non-detections). 518 of the calculations (25.2%) have ratios below 2 (i.e. the target maximum is also the global maximum, but the detection may require a threshold that also results in an unacceptable number of false alarms). For detecting the S-waves using the C trace, there is no change in the percentage of strict non-detections. However, 1766 out of 2550 calculations (69.25%) fail to result in a ratio of target maximum to highest noise value above 2. There are no hard and fast rules as to where we should choose a cut-off value for selecting a detection threshold but, for this particular set of signal detection challenges, it would seem that the use of the $C|C|$ trace provides a real improvement over C traces for a separation between the values associated with the target signals and the values associated with noise.

The experiment described here clearly uses an idealized set of data: short data segments with exactly one positive match in each window and an absence of signals from other events which greatly resemble the target signals. The signals are also of fairly similar amplitude and signal-to-noise ratio. In a situation of monitoring natural or induced seismicity we would anticipate both a wide range of signal amplitude and SNR and a slow diminishing of waveform similarity as hypocentral distance and/or source mechanism difference increases. The short template length was also selected to be "marginal", i.e. such that a significant number of signals failed to be found. Had we used a template length of 6s (as in Gibbons et al. 2020) then all signals would have been detected confidently and the highest noise values would have been reduced significantly. What we have done is ultimately to demonstrate that the $C|C|$ trace has a greater contrast between values at times of interest and values at other times. In the next section, we seek to provide confidence that there is a basis for choosing the $C|C|$ trace that makes it preferable to some other transformation that also stretches out the significant values relative to the background values.

4   THE CONNECTION BETWEEN CORRELATION AND REGRESSION

An almost arbitrarily chosen statistics text book will describe how the square of the correlation co-efficient, sometimes referred to as the Coefficient of Determination, C^2, indicates the fraction of the variance of one variable that can be explained by another variable in a regression. This is exactly what

we seek in our pattern detector: a function that returns a high value when a window $\boldsymbol{y}_k$ of the incoming data is well explained by the waveform template $\boldsymbol{x}$ and a low value when it is not well explained. In a linear regression, we seek scalars α_k and β_k such that the residual vector

$$\boldsymbol{\varepsilon}_k = \boldsymbol{y}_k - \alpha_k \boldsymbol{x} - \beta_k \tag{5}$$

is a minimum. A few lines of code will confirm for successive windows $\boldsymbol{y}_k$ of our incoming data that

$$C_k^2 \approx \frac{|\boldsymbol{y}_k|^2 - |\boldsymbol{\varepsilon}_k|^2}{|\boldsymbol{y}_k|^2} \tag{6}$$

for the optimal α_k and β_k, where C_k^2 is the square of our fully normalized correlation coefficient for this data window. In principle, we could implement a detector which performed a linear regression on subsequent data windows and attempted to measure some distance metric between the template vector and the vector of incoming data. Given the relationship in Eq. 6 we would not of course do this explicitly since C_k^2 is calculated very rapidly in the frequency domain. To see what different values of our two detection statistics mean in terms of the ratio

$$R_k = \frac{|\boldsymbol{\varepsilon}_k|}{|\boldsymbol{y}_k|} \tag{7}$$

we simply rearrange Eq. 6 to give

$$R_k = \sqrt{1 - C_k^2} \tag{8}$$

which we could denote the "unexplained fraction". We can similarly denote $1 - R$ as the "explained fraction" or signal fraction. This is not to say that there is a signal present in this window of the incoming data; this is just the fraction of this data window which could be explained by the presence of a signal the same shape as the waveform template. We plot the explained fraction in Fig. 5, both as a function of our two detection statistics, C and $C|C|$, and an alternative candidate function, C^4. C^4 would lead to an even more non-linear transformation of the correlation trace. While such a transformation would increase the ratio between values corresponding to good signal matches and values corresponding to unrelated noise, it would also risk pushing more values into the tail of the distribution and raising the false alarm rate.

What we want in a matched signal detector is a diagnostic which differentiates as confidently as possible between a window of data (potentially) containing a signal resembling our template and a window of data which does not contain such a signal. Imagine two points on Fig. 5, one for a data window well explained by the template waveform and one for a data window poorly explained by this template. The ratio between the y-coordinates of these points measures the contrast between how well the two vectors match the template. The ratio between the x-coordinates of these points measures how easily we would detect the signal against the background noise. If we reduce $C|C|$ by a given factor, we reduce the signal fraction by almost the same factor. If we reduce C by a given factor, we

reduce the signal fraction by a considerably higher factor. Equivalently, the SNR for C increases more slowly than the SNR for $C|C|$ as the fit to the template improves. As for the C^4 detection statistic, the non-linearity results in a rapid decrease in the fit to the template over a relatively narrow range of the detection statistic.

5   CONCLUSIONS AND RECOMMENDATIONS

There are many variants of the matched signal detector in operational use in contemporary seismology. Where applicable, they can provide an exceptionally low detection threshold. We need however to be constantly aware of the limitations of correlation detectors, realizing how rapidly waveform similarity diminishes as the distance between hypocenters increases (e.g. Menke 1999; Harris & Dodge 2020). (Gao et al. 2021, also point out that deceptively similar waveforms often occur for events that are not closely located.) There are numerous strategies for detection and/or identification of seismic events where the form of the signal varies from the template in one way or another (e.g. Harris & Kværna 2010; Carmichael 2016). There are also numerous strategies for mitigating performance degradation due to challenges such as noise transients and challenging amplitude characteristics of signals (e.g. Gibbons et al. 2012; Gao & Kao 2020).

The toolbox available for autonomous detection, location, and classification of seismic events in near real-time is expanding rapidly, with machine learning likely to take a dominant role (e.g. Shen & Shen 2021). Auto-correlation detectors, searching for repeating events without prior templates, (e.g. Brown et al. 2008) are essentially performing unsupervised learning. The FAST (Fingerprint And Similarity Threshold) algorithm (Yoon et al. 2015) increases general applicability and computational efficiency without sacrificing signal detectability. Machine learning is likely to both complement and partially replace the role of correlation detectors. Even without considering machine learning, correlation detectors and related methods are going to represent only a part of the solution for the monitoring of low magnitude seismic events (e.g. Li et al. 2020; Ichinose et al. 2021). However, there is no more sensitive detector than a correlator for a situation where we know the exact form of a signal.

One aspect of the seismic correlation detector which has arguably not been addressed with the rigour it deserves is the form of the detection statistic itself. Some procedures calculate the fully normalized correlation coefficient, C. Others use, either implicitly or explicitly, the square of the correlation coefficient, C^2. C is appealing for use in multi-channel correlation detectors because of how classical beamforming methods can be applied to these traces from different channels: both to increase SNR and to identify false alarms by checking the alignment. An alternative detection statistic, $C|C|$, has the dynamic range of C^2 but varies in sign like C. We demonstrate in this paper that there are differences between the performance of a detector that calculates C and a detector that calculates

$C|C|$; the SNR for C increases more slowly than for $C|C|$ as the fit to the applied waveform template improves. We demonstrate this for a convenient but somewhat limited set of seismic signals on 3-component data.

It is unclear as to how significant this difference will become on larger-scale operational correlation detection systems. As heavy stacking of traces over networks suppresses the background level of the detection traces, the differences may become academic. This will need to be investigated in controlled experiments using tried and tested software on diverse sets of seismicity. Although the dataset used in the present study has multiple event pairs with differing degrees of waveform similarity, it will be necessary to test how the performance of the different detection statistics changes with factors such as signal amplitude, event magnitude, source mechanism, source-time history, and network configuration. We have seen one example where a spurious detection is visually more convincing with the $C|C|$ detection statistic than with C (Figure 2). However, the subsequent investigation performed does not indicate a decrease in performance and we recall that the alignment post-processing procedure for false-alarm screening (Gibbons & Ringdal 2006) can be applied to $C|C|$ as readily as with C. There will likely be increasing efforts to compare the performance of different matched signal detectors (e.g. De La Hoz et al. 2021) and it is important that we compare like with like and that we understand and document the performance of all algorithms compared. There is likely much to be gained by visualizing the distribution of detection statistics in operational pipelines, both for quality control and assessment of detection sensitivity.

There are many metrics by which the performance of a matched signal detection system can be evaluated. Computational efficiency aside, the target is generally a Receiver Operating Characteristic (ROC)-curve that optimizes the ratio between the True Positive Rate and the False Positive Rate. The definitions of these quantities are determined by user requirements. Are we seeking to characterize all seismicity in a given region or are we seeking signals from one very specific source, treating all other signals as noise? The recipe proposed by Gibbons & Ringdal (2006) for detecting near-repeating seismic events on multi-channel waveform data is both powerful, robust, and widely applicable. However, the detection statistic trace, C, that paper recommended is not the optimal one and performance of operational correlation detection frameworks can probably be improved by instead calculating a $C|C|$ trace.

ACKNOWLEDGMENTS

gust 2021) and the free GMT software was obtained from `https://www.generic-mapping-tools.org/` (last accessed August 2021).

REFERENCES

Abbott, B. P. et al., 2016, GW150914: First results from the search for binary black hole coalescence with Advanced LIGO, *Physical Review D*, **93**(12), 122003.

Albuquerque Seismological Laboratory (ASL)/USGS, 1988, Global Seismograph Network (GSN - IRIS/USGS).

Bachura, M. & Fischer, T., 2019, Waveform CrossCorrelation for Differential Time Measurement: Bias and Limitations, *Seismological Research Letters*, **90**(5), 2005–2014.

Barrett, S. A. & Beroza, G. C., 2014, An Empirical Approach to Subspace Detection, *Seismological Research Letters*, **85**(3), 594–600.

Beaucé, E., Frank, W. B., & Romanenko, A., 2018, Fast Matched Filter (FMF): An Efficient Seismic Matched-Filter Search for Both CPU and GPU Architectures, *Seismological Research Letters*, **89**(1), 165–172.

Benz, H. M., McMahon, N. D., Aster, R. C., McNamara, D. E., & Harris, D. B., 2015, Hundreds of earthquakes per day: The 2014 Guthrie, Oklahoma, earthquake sequence, *Seismological Research Letters*, **86**(5), 1318–1325.

Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J., 2010, ObsPy: A Python Toolbox for Seismology, *Seismological Research Letters*, **81**(3), 530–533.

Brown, J. R., Beroza, G. C., & Shelly, D. R., 2008, An autocorrelation method to detect low frequency earthquakes within tremor, *Geophysical Research Letters*, **35**(16), L16305.

Brunelli, R., 2009, *Template Matching Techniques in Computer Vision*, John Wiley and Sons Ltd, Chichester, United Kingdom, doi:10.1002/9780470744055, ISBN 978-0-470-51706-2, 348 pp.

Carmichael, J. D., 2016, A Waveform Detector That Targets Template–Decorrelated Signals and Achieves Its Predicted Performance, Part I: Demonstration with IMS Data, *Bulletin of the Seismological Society of America*, **106**(5), 1998–2012.

Carmichael, J. D. & Hartse, H., 2016, Threshold Magnitudes for a Multichannel Correlation Detector in Background Seismicity, *Bulletin of the Seismological Society of America*, **106**(2), 478–498.

Chamberlain, C. J., Hopp, C. J., Boese, C. M., WarrenSmith, E., Chambers, D., Chu, S. X., Michailos, K., &

Townend, J., 2018, EQcorrscan: Repeating and NearRepeating Earthquake Detection and Analysis in Python, *Seismological Research Letters*, **89**, 173–181.

Chamberlain, C. J., Townend, J., & Gerstenberger, M. C., 2020, RT-EQcorrscan: Near-Real-Time Matched-Filtering for Rapid Development of Dense Earthquake Catalogs, *Seismological Research Letters*, **91**.

De La Hoz, C., Tary, J. B., & Lomax, A., 2021, Empirical subspace detection applied to triggered seismicity by the July 25, 2011, M 5.0 earthquake in the Sea of Marmara, Turkey, *Computers & Geosciences*, **152**(April), 104738.

Duboeuf, L., Oye, V., & Dando, B. D. E., 2021, An Aggregated Template (AT) methodology: novel automatic phaseonset identification by template matching, *Geophysical Prospecting*.

Ford, S. R. & Walter, W. R., 2015, International Monitoring System Correlation Detection at the North Korean Nuclear Test Site at Punggye-ri with Insights from the Source Physics Experiment, *Seismological Research Letters*, **86**, 1160–1170.

Gao, D. & Kao, H., 2020, Optimization of the Match-Filtering Method for Robust Repeating Earthquake Detection: The Multisegment Cross-Correlation Approach, *Journal of Geophysical Research: Solid Earth*, **125**(7), 1–19.

Gao, D., Kao, H., & Wang, B., 2021, Misconception of Waveform Similarity in the Identification of Repeating Earthquakes, *Geophysical Research Letters*, **48**(13), 1–12.

Geller, R. J. & Mueller, C. S., 1980, Four similar earthquakes in central California, *Geophysical Research Letters*, **7**(10), 821–824.

Gibbons, S., Näsholm, S. P., & Kværna, T., 2014, A multi-station matched filter and coherent network processing approach to the automatic detection and relative location of seismic events, doi:10.6084/m9.figshare.1226559.v1, Geophysical Research Abstracts, Vol. 16, EGU2014-11880, EGU General Assembly 2014.

Gibbons, S. J. & Ringdal, F., 2006, The detection of low magnitude seismic events using array-based waveform correlation, *Geophys. J. Int.*, **165**, 149–166.

Gibbons, S. J. & Ringdal, F., 2012, Seismic Monitoring of the North Korea Nuclear Test Site Using a Multi-channel Correlation Detector, *Geoscience and Remote Sensing, IEEE Transactions on*, **50**, 1897–1909.

Gibbons, S. J., Ringdal, F., & Kværna, T., 2012, Ratio-to-moving-average seismograms: A strategy for improving correlation detector performance, *Geophysical Journal International*, **190**(1), 511–521.

Gibbons, S. J., Kværna, T., Näsholm, S. P., & Mykkeltveit, S., 2018, Probing the DPRK Nuclear Test Site down to Low Seismic Magnitude, *Seismological Research Letters*, **89**(6), 2034–2041.

Gibbons, S. J., Kværna, T., Tiira, T., & Kozlovskaya, E., 2020, A Benchmark Case Study for Seismic Event Relative Location, *Geophysical Journal International*, **223**(2), 1313–1326.

Harris, C. R. et al., 2020, Array programming with NumPy, *Nature*, **585**(7825), 357–362.

Harris, D., Gibbons, S., Rodgers, A., & Pasyanos, M., 2012, Nuclear Test Ban Treaty Verification: Improving Test Ban Monitoring with Empirical and Model-Based Signal Processing, *IEEE Signal Processing Magazine*, **29**(3), 57–70.

Harris, D. B., 2006, Subspace detectors: theory, Tech. Rep. UCRL-TR-222758, Lawrence Livermore National Laboratory, Livermore, California.

Harris, D. B. & Dodge, D. A., 2011, An Autonomous System for Grouping Events in a Developing Aftershock Sequence, *Bulletin of the Seismological Society of America*, **101**(2), 763–774.

Harris, D. B. & Dodge, D. A., 2020, The Geometry of Signal Space: A case study of direct mapping between seismic signals and event distribution, *Geophysical Journal International*, pp. 2189–2208.

Harris, D. B. & Kværna, T., 2010, Superresolution with Seismic Arrays using Empirical Matched Field Processing, *Geophys. J. Int.*, **182**, 1455–1477, doi:10.1111/j.1365-246X.2010.04684.x.

Harris, D. B. & Paik, T., 2006, Subspace Detectors: Efficient Implementation, Tech. Rep. UCRL-TR-223177, Livermore, CA.

Hicks, S. P., Verdon, J., Baptie, B., Luckett, R., Mildon, Z. K., & Gernon, T., 2019, A Shallow Earthquake Swarm Close to Hydrocarbon Activities: Discriminating between Natural and Induced Causes for the 20182019 Surrey, United Kingdom, Earthquake Sequence, *Seismological Research Letters*, **90**, 2095–2110.

Hughes, L., Chamberlain, C. J., Townend, J., & Thomas, A. M., 2021, A Repeating Earthquake Catalog From 2003 to 2020 for the Raukumara Peninsula, Northern Hikurangi Subduction Margin, New Zealand, *Geochemistry, Geophysics, Geosystems*, **22**(5), 1–20.

Ichinose, G. A., Ford, S. R., Kroll, K., Dodge, D., Pyle, M., Pitarka, A., & Walter, W. R., 2021, Preliminary Analysis of Source Physics Experiment ExplosionTriggered Microseismicity using the BackProjection Method, *Journal of Geophysical Research: Solid Earth*.

Junek, W. N., Kværna, T., Pirli, M., Schweitzer, J., Harris, D. B., Dodge, D. A., & Woods, M. T., 2015, Inferring Aftershock Sequence Properties and Tectonic Structure Using Empirical Signal Detectors, *Pure and Applied Geophysics*, **172**(2), 359–373.

Li, L., Tan, J., Schwarz, B., Staněk, F., Poiata, N., Shi, P., Diekmann, L., Eisner, L., & Gajewski, D., 2020, Recent Advances and Challenges of Waveform-Based Seismic Location Methods at Multiple Scales, *Reviews of Geophysics*, **58**(1), 1–47.

Menke, W., 1999, Using waveform similarity to constrain earthquake locations, *Bulletin of the Seismological Society of America*, **89**(4), 1143–1146.

Parvulescu, A., 1995, Matchedsignal (MESS") processing by the ocean, *The Journal of the Acoustical Society of America*, **98**(2), 943–960.

Pirli, M., Gibbons, S. J., & Schweitzer, J., 2011, Application of array-based waveform cross-correlation techniques to aftershock sequences: The 2003 Lefkada Island, Greece, case, *Journal of Seismology*, **15**(3), 533–544.

Richards, P., Waldhauser, F., Schaff, D., & Kim, W.-Y., 2006, The Applicability of Modern Methods of Earthquake Location, *Pure and Applied Geophysics*, **163**(2), 351–372.

Schaff, D. P., 2008, Semiempirical Statistics of Correlation-Detector Performance, *Bulletin of the Seismological Society of America*, **98**(3), 1495–1507.

Shearer, P. M., 1997, Improving local earthquake locations using the L1 norm and waveform cross correlation:

Application to the Whittier Narrows, California, aftershock sequence, *J. Geophys. Res.*, **102**(B4), 8269–8283.

Shelly, D. R., Beroza, G. C., & Ide, S., 2007, Non-volcanic tremor and low-frequency earthquake swarms, *Nature*, **446**, 305–307.

Shen, H. & Shen, Y., 2021, Array-Based Convolutional Neural Networks for Automatic Detection and 4D Localization of Earthquakes in Hawaii, *Seismological Research Letters*.

Slinkard, M., Schaff, D., Mikhailova, N., Heck, S., Young, C., & Richards, P. G., 2014, Multistation Validation of Waveform Correlation Techniques as Applied to Broad Regional Monitoring, *Bulletin of the Seismological Society of America*, **104**(6), 2768–2781.

Slinkard, M. E., Carr, D. B., & Young, C. J., 2013, Applying Waveform Correlation to Three Aftershock Sequences, *Bulletin of the Seismological Society of America*, **103**(2A), 675–693.

Stamoulis, C. & Richardson, A. G., 2010, Application of matched filtering to identify behavioral modulation of brain oscillations, *Journal of Computational Neuroscience*, **29**(1-2), 63–72.

Stevens, J. L., Gibbons, S., Rimer, N., Xu, H., Lindholm, C., Ringdal, F., Kvaerna, T., & Murphy, J. R., 2006, Analysis and simulation of chemical explosions in nonspherical cavities in granite, *Journal of Geophysical Research*, **111**(B4), B04306.

Taylor, S. R., Arrowsmith, S. J., & Anderson, D. N., 2013, Development of a matched filter detector for acoustic signals at local distances from small explosions, *The Journal of the Acoustical Society of America*, **134**(1), EL84–EL90.

Turin, G., 1960, An introduction to matched filters, *IEEE Transactions on Information Theory*, **6**(3), 311–329.

van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., & Yu, T., 2014, scikit-image: image processing in Python, *PeerJ*, **2**(1), e453.

van Trees, H. L., 1968, *Detection, Estimation and Modulation Theory: Vol I*, John Wiley and Sons, New York, USA.

van Wijk, K., Chamberlain, C. J., Lecocq, T., & Van Noten, K., 2021, Seismic monitoring of the auckland volcanic field during new zealand's covid-19 lockdown, *Solid Earth*, **12**(2), 363–373.

VanDecar, J. C. & Crosson, R. S., 1990, Determination of Teleseismic Relative Phase Arrival Times Using Multi-Channel Cross-Correlation and Least Squares, *Bulletin of the Seismological Society of America*, **80**, 150–169.

Warren-Smith, E., Fry, B., Kaneko, Y., & Chamberlain, C. J., 2018, Foreshocks and delayed triggering of the 2016 MW7.1 Te Araroa earthquake and dynamic reinvigoration of its aftershock sequence by the MW7.8 Kaikura earthquake, New Zealand - ScienceDirect, *Earth and Planetary Science Letters*, **482**, 265–276.

Wessel, P., Luis, J. F., Uieda, L., Scharroo, R., Wobbe, F., Smith, W. H. F., & Tian, D., 2019, The Generic Mapping Tools Version 6, *Geochemistry, Geophysics, Geosystems*, **20**(11), 5556–5564.

Yoon, C. E., O'Reilly, O., Bergen, K. J., & Beroza, G. C., 2015, Earthquake detection through computationally efficient similarity search, *Science Advances*, **1**(11), e1501057.

Zhang, M. & Wen, L., 2015, An effective method for small event detection: match and locate (M&L), *Geophysical Journal International*, **200**, 1523–1537.
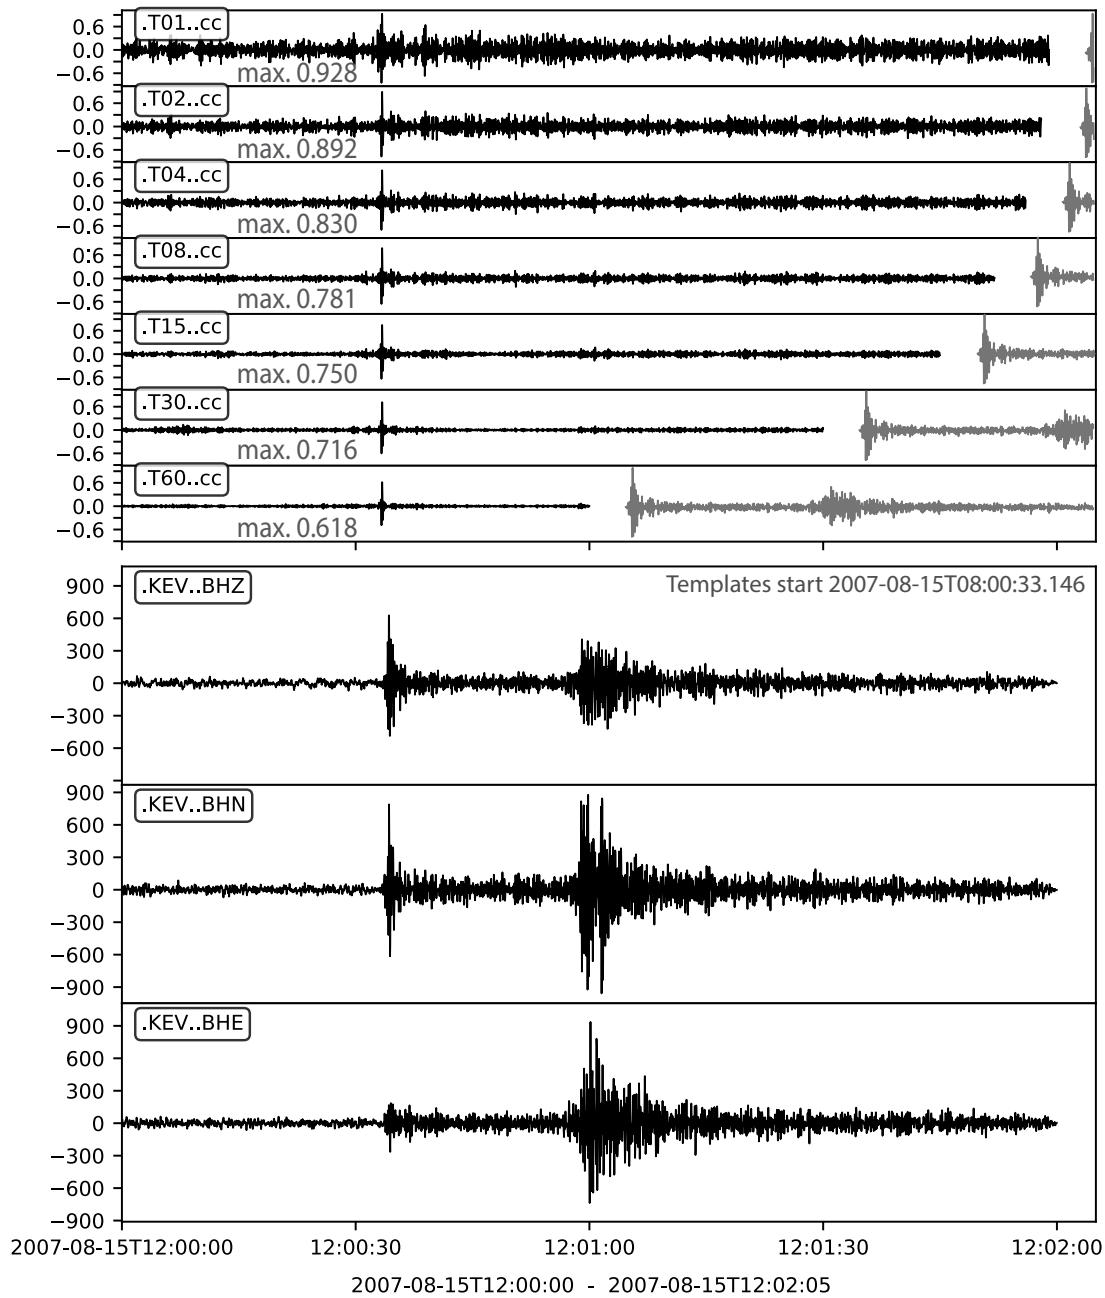
Figure 1. Signal filtered 2-8 Hz (lowermost 3 traces) on the 3 components of station IU.KEV of event H02 from Gibbons et al. (2020) and correlation coefficient traces $C(t)$ (3-component mean) obtained from correlating against the corresponding templates of different lengths from event H01 (vertical component only displayed in grey alongside $C(t)$ traces displayed in black).
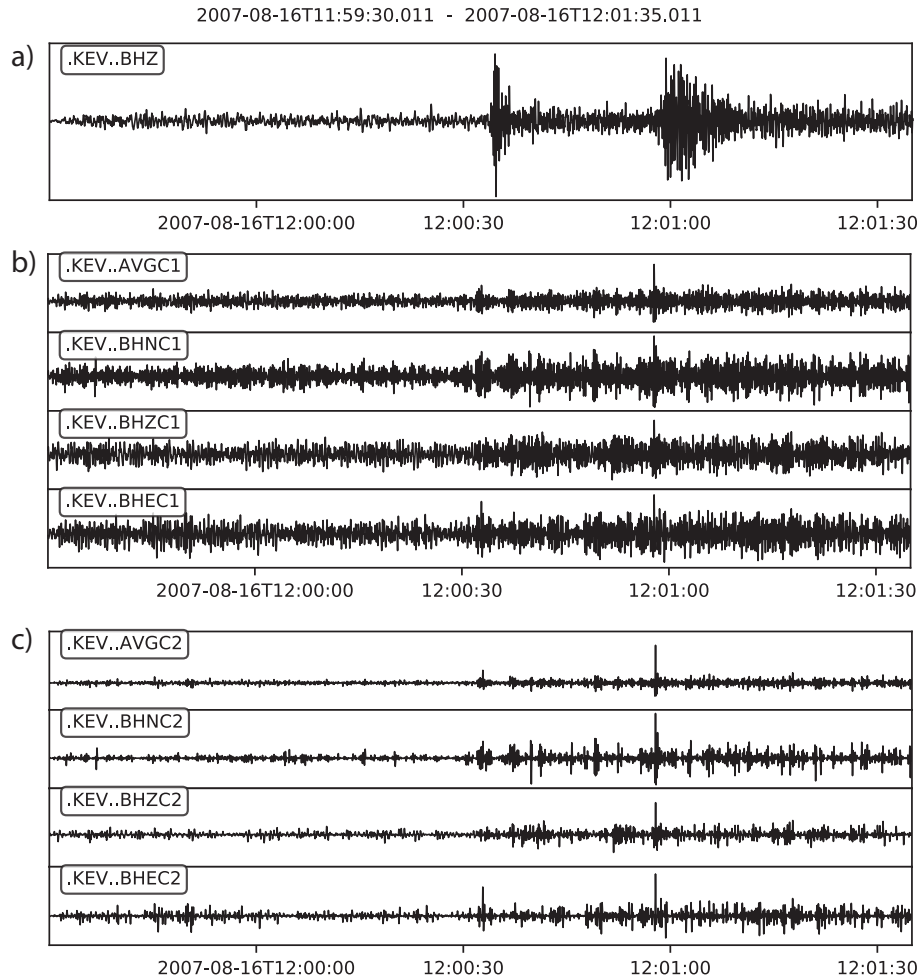
Figure 2. Detection of the S-phase from event H04 in Gibbons et al. (2020) using a 2s long waveform template from event H03 starting at a time 2007-08-16T08:00:55.713. (a) displays the vertical component waveform for event H04 (bandpass filtered 2-8 Hz), (b) displays the single-channel and multi-channel correlation coefficient traces, C, (labelled C1) and (c) displays the corresponding detection traces but with C replaced with $C|C|$ (labelled C2). In panel (b), the trace labelled AVGC1 is the sample-wise mean of the three single-channel correlation (C) traces below: BHNC1, BHNC1, and BHEC1. In panel (c), the trace labelled AVGC2 is the sample-wise mean of the three single-channel $C|C|$ traces below: BHNC2, BHNC2, and BHEC2.
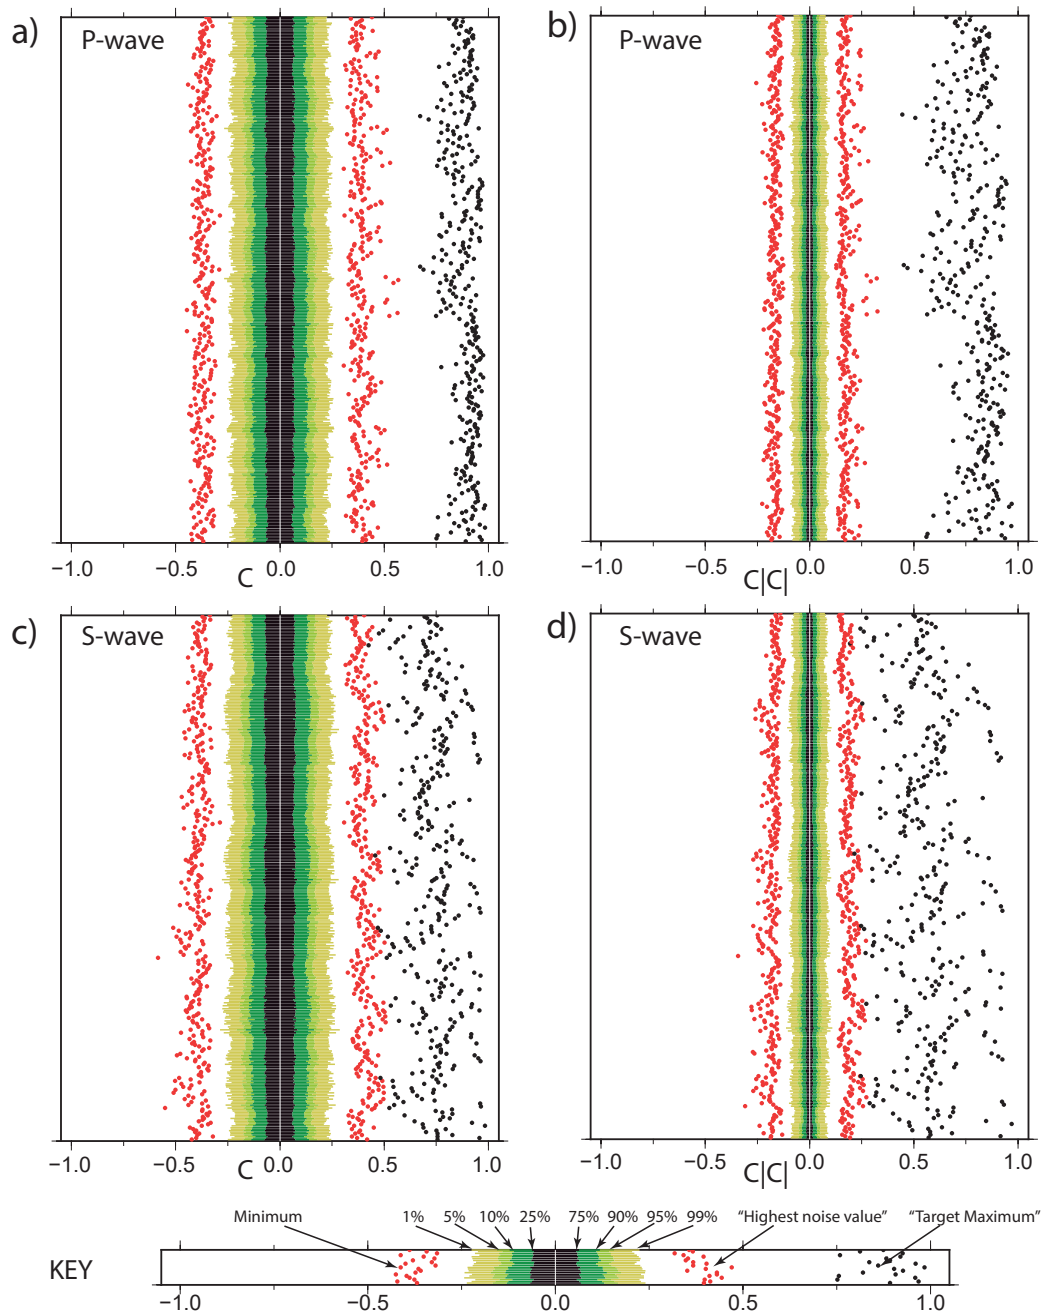
Figure 3. Distributions of values on 3-component average C and $C|C|$ traces when attempting to detect P- and S-waves from 51 explosions in the Gibbons et al. (2020) dataset using only the IU.KEV waveforms provided. Each line displays the minimum and maximum values for a single template detection with percentiles as displayed in the KEY below. Templates of length 2s are used starting at the times prescribed in Gibbons et al. (2020). Of a total of 2550 event pairs, the first 350 are displayed in each panel. The target maximum value is the maximum value within 2.5s of the anticipated arrival time and the highest noise value is the maximum value of the set of samples that are further than 2.5s from the predicted arrival time.
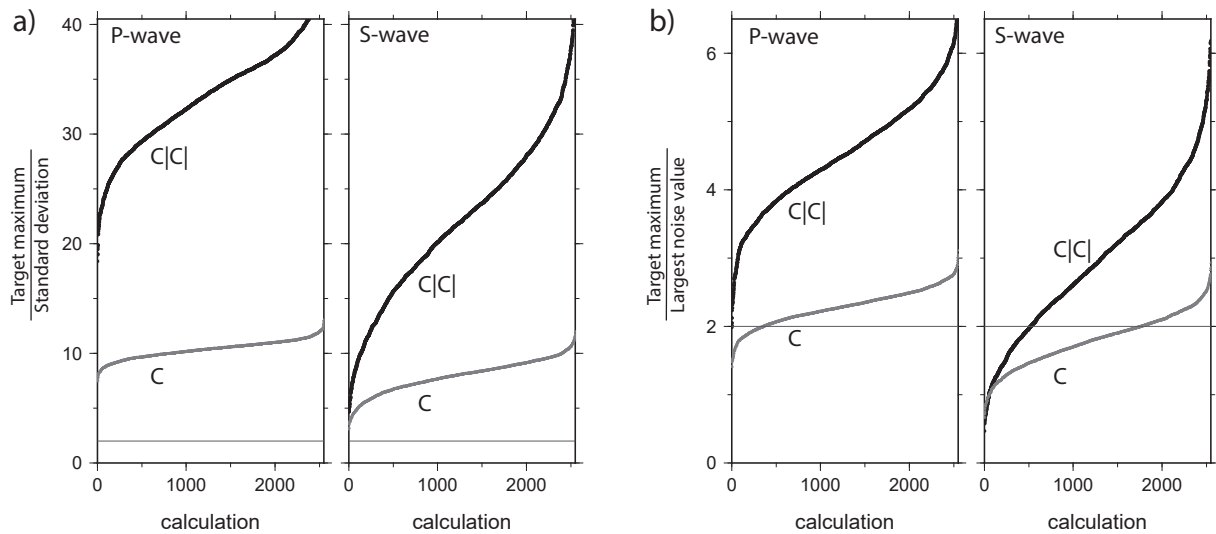
Figure 4. Comparison of key characteristics of 3-component average C and $C|C|$ traces for all 2550 event pairs. Panel a) shows the ratio between the target maximum value and the standard deviation and panel b) shows the ratio between the target maximum value and the "Maximum noise value". The target maximum value is the maximum value within 2.5s of the anticipated arrival time and the highest noise value is the maximum value of the set of samples that are further than 2.5s from the predicted arrival time. The $y = 2$ line is an arbitrarily chosen ratio threshold and is just displayed as potential threshold between "detected" and "not detected". Each line is really a scatter-plot of points, ordered from smallest value of the ratio displayed to the largest. The event pair that each point plotted corresponds to therefore varies from line to line.
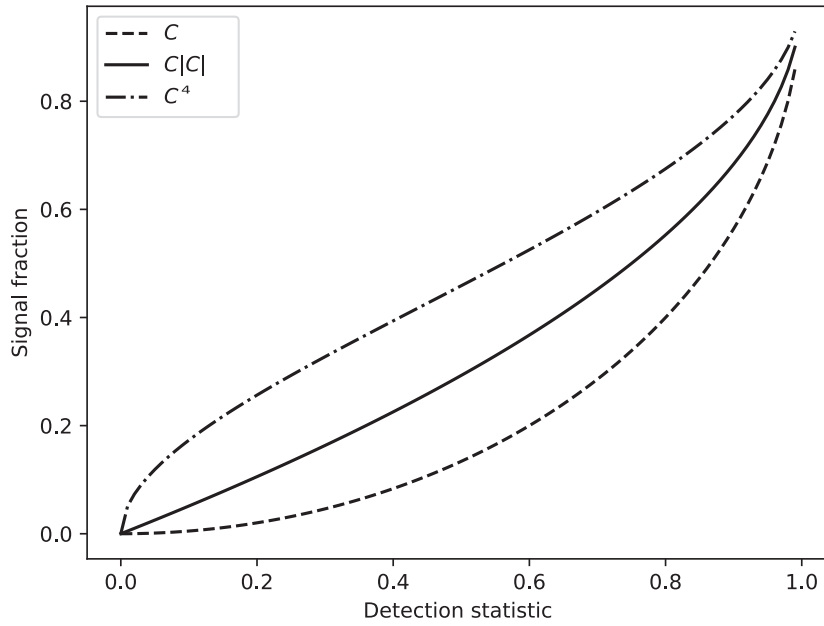
Figure 5. If R denotes the ratio $|\boldsymbol{\varepsilon}| : |\boldsymbol{y}|$, where $\boldsymbol{y}$ is a vector of the incoming data and $\boldsymbol{\varepsilon}$ is the residual vector resulting from a linear regression with the template vector $\boldsymbol{x}$, then $R \approx \sqrt{1 - C^2}$ where C is the fully normalized correlation coefficient. We denote $(1 - R)$ the signal fraction and display this as a function of detection statistic for three candidates: C, $C|C|$, and C^4. No data is considered for this plot, only the assumed relation $R = \sqrt{1 - C^2}$. The quantity $1 - \sqrt{1 - C^2}$ is more precisely described as the *explained fraction*, since the term *signal fraction* implies that the remainder is noise: see text for details.