

Using Approximate Nearest Neighbor Methods to Rapidly Search Large Seismic Signal Archives

Chris Young & Antonio Gonzales

Sandia National Laboratories

Motivation (1)

- **Waveform correlation has proven to be a very effective means to detect and identify signals from various types of events...**

Bulletin of the Seismological Society of America, Vol. 81, No. 6, pp. 2395–2410, December 1991

A WAVEFORM CORRELATION METHOD FOR IDENTIFYING QUARRY EXPLOSIONS

By D. B. HARRIS

ABSTRACT

A waveform correlation method is presented for identifying quarry explosions by attributing them to known mines characterized by multiple master events. The objective is to provide a reliable automatic procedure for screening the large number of quarry explosions likely to be detected by networks of in-country stations, including a balanced set of test-bed variables. The method generalizes existing correlation techniques to compare waveform from an unknown event recorded at an array of sensors with a linear combination of master event waveforms recorded at the same array. The use of a linear combination reduces the chance of a mismatch between the waveforms and minimizes the chance of spectral excitation between the events being compared. The weights in the linear combination are filters, offering some compensation for variations in source time functions and errors of waveform alignment. The use of array data reduces the chance of false attribution by reducing bias and variance in the correlation measurement.

In a test conducted with P-wave data segments recorded at a 13-element array, the method correctly identified 100% of the quarry events. The method is marginal due to waveform correlation noise. The correlation waveform detection threshold is set at 0.7 of the smaller TB with the related uncertainty. The events are again events in the approximation, the second event from events 4 array its resolution.

Bulletin of the Seismological Society of America, Vol. 80, No. 6, pp. 2177–2193, December 1990

CORRELATION OF WAVEFORMS FROM CLOSELY SPACED REGIONAL EVENTS

By HANS ISRAELSON

ABSTRACT

We studied the similarity of waveforms recorded at the high frequency element of the NORESS array from 137 events (M = 0.8 to 2.1) in or near the mining districts of Central Sweden. Waveform correlation based on the covariance matrices of three-component recordings on one minimum amplitude calculated from three filtered segments of 0.5 and 1.5 s in length and the signal-to-noise ratio consistently peaked for the waveforms. A cutoff value for the correlations to regional waveforms with poor and good correlation values could be determined from the statistical uncertainty of the correlation of the NORESS. More than 80 per cent of the events could be grouped with hierarchical clustering of one large group of 98 events and five smaller groups. The NORESS epicenters of the events in the large group were scattered over an area of 20 × 75 km. An expansion of the data set to 130 events from the same region and with a larger TB could be fitted to the data of this group (exp. ~ 7), taking NORESS epicenter uncertainties into account. A procedure for location of close events based on correlation values is described. The events in the large event group fit in the events recorded in an area of 40 × 100 km. The five smaller groups high correlation values were obtained above 15 Hz. These frequencies are significantly higher than those reported in other studies, and the so-called quarter wavelength argument constrains the event epicenters to within 0.1 km.

INTRODUCTION

In regional seismic monitoring one sometimes distinguishes between model-based and case-based approaches at defining seismic events (Dysart and Puli, 1987). model-based approach aims at defining some discriminant from the recorded waveform that can in general be related to the source. The discriminant should include the physics of various seismic source types, such as the ratio of radiated P - and S -wave energy. In spite of a number of studies, no single model-based method has so far been defined that performs successfully regardless of region. In contrast, the case-based approach attempts to take advantage of the repeatability of recorded events and same source types for location and identification. Similarities of waveform recordings from closely spaced sources have been reported for several regions in the literature. For example, seismograms recorded at the NORESS array from several events at the Tittan mine (400 km distance) showed significant similarity in time history (Jukarveki, 1987). Records may not only be recorded at a particular location, but also at other locations, such as nearby locations. Many differences have, for example, been observed for events at NORESS from events in the Western U.S.S.R. at neighboring mines only 25 km apart (Suteau-Henson and Bache, 1988). From visual inspection of recordings of mining explosions, experienced analysts have even been able to pinpoint the particular mine of origin. The case-based approach is thus based on knowledge acquired from the recordings of many events from a particular area, and a new event is supposed to be identified by waveform comparisons with events previously recorded in the region.

2177

JOURNAL OF GEOPHYSICAL RESEARCH, VOL. 116, B03309, doi:10.1029/2010JB008095, 2011

On finding and using repeating seismic events in and near China

David P. Schaff and Paul G. Richards¹

Received 9 August 2010; revised 11 December 2010; accepted 3 January 2011; published 23 March 2011.

[1] A comprehensive search for repeating seismic events was performed using cross correlation on waveforms for 17,898 seismic events in and near China reported in the Annual Catalog of Chinese Earthquakes from 1985 to 2005. We define a repeat as an event pair having a time cross correlation of 0.4 for a particular time interval (0.5 s) and time window (5 s before P arrival to 40 s after Lg arrival). Such events typically have hypocenters separated by no more than 1 km. This result enables a determination of median event repeat times and standard deviations for each event type that locate events in this study one at a time using phase picks. The maximum location error is on the order of hundreds of kilometers. We use normalized cross-correlation values to measure relative amplitude of event pairs. The standard deviation is about one magnitude unit below that of the maximum magnitude for doublet events, which we estimate for the Chinese catalog to have a standard deviation of 0.36 magnitude units. Two thousand three hundred and seventy-nine events of the 17,898 events (13%) in the events in this catalog are classified as repeats, with later years showing an increase in the percentage. About half the events occur as isolated doubles, the rest as multiples of three. The events appear to be triggered or are due to natural processes. The results show that

Chadwick, Schaff, P. G. Richards, *J. Geophys. Res.*, 116, B03309, doi:10.1029/2010JB008095.

Geophys. J. Int. (2006) 165, 149–166

The detection of low magnitude seismic events using waveform correlation

Steven J. Gibbons and Frode Ringdal

NOAA, PO Box 35, 2027 Kalaeloa, Hawaii. E-mail: sjgibbons@noaa.gov

Accepted 2005 November 6; Received 2005 September 7, in original form 2005 June 11

SUMMARY

It has long been accepted that occurrences of a known signal cross-correlating the incoming data stream with a waveform detector have received very little attention in the field of an automatic seismic event detection system. We present a method for detecting repeating events in highly correlated geophysical signals. Repeating events in highly correlated geophysical signals are similar waveforms and good signals from events at a given location. They are detected at lower magnitudes than by power detectors. Even greater improvements in signal detection arrays, running correlation coefficients from single sensor network to result in a network correlation coefficient that two events are correlated. The detection threshold is indicated by the cross-correlation function is identical for all means that the correlation coefficient traces are coherent. We illustrate the method using data from the 2004 Sumatra Sea event. The weak event that occurred 4 h after the main STA detector on the SPTS array but is readily detected channel by channel. The STA detector is the standard conventional detection can be made for the second event. When the correlation coefficients are beamformed over a reduced time window, the detection threshold of the cross-correlation by scaling down a master signal and in imaging for one case, waveform correlation using a single channel detects signals of approximately 0.7 orders of magnitude lower than possible using a STA detector on a array beam. Waveform correlation on a array provides a significant improvement in detection threshold for low magnitude events.

We describe a case study in which small seismic events at the Barentsburg coal mine on Spitsbergen were detected using the signals from a major rockburst as master waveforms. The events were detected using the STA detector, which exhibited coincidental similarity with unrelated incoming waveforms. We demonstrate how such false alarms can almost always be identified and screened out automatically by performing frequency-wavenumber analysis upon the set of individual correlation coefficient traces.

Key words: detection threshold, matched signal, seismic array, waveform correlation

1 INTRODUCTION

The most effective method of detecting a known signal in a potentially noisy time-series is to cross-correlate a waveform template with successive time segments of incoming data. Any segments of the incoming data that are sufficiently similar to the template or master waveform will result in a high value of the correlation function. Such a procedure is referred to as a matched filter or a matched signal detector (see e.g. van der Veen 1983). The

detection by waveform correlation of a synthetic signal buried in noise was demonstrated by Asteny (1968) in a paper that provided the first quantitative analysis of the use of waveform correlation in signal recognition in exploration geophysics and other fields.

Such methods have so far received very little attention in the field of seismic event detection. The STA detector is the standard detector to the master or template waveform and the over-whelming majority of signals detected by the STA detector are subspace detectors come from inhomogeneous sources and consequently have

© 2006 The Authors
Journal compilation © 2006 RAS

Bulletin of the Seismological Society of America, Vol. 104, No. 6, pp. 2768–2781, December 2014, doi:10.1785/0120140104

Multistation Validation of Waveform Correlation Techniques as Applied to Broad Regional Monitoring

by Megan Slinkard, David Schaff, Natalya Mikhalkova, Stephen Heck, Christopher Young, and Paul G. Richards

Abstract Waveform correlation is garnering attention as a method for detecting, locating, and determining the magnitude of seismic events. In addition, due to its potential for use in broad regional monitoring, we applied the technique to a large region of central Asia over a three-year period, monitoring for events at regional distances using three high-quality stations. We discuss methods for choosing quality templates and introduce a method for choosing correlation detection thresholds, tailored for each template, for a desired false alarm rate. Our *SeisCorr* software found more than 10,000 events in the catalog. We validate the detections using three methods: cross-correlation with a template, multistation validation, and correlation with a doublet. We also validate detections using three methods: cross-correlation with a template, high correlation with a doublet, and multistation validation. At each station, 65%–75% of our detections could be confirmed, most by multistation validation. We confirmed over 6500 unique detections. From these observations, it is of interest that a significant portion of the Comprehensive Nuclear-Test-Ban Treaty Organization's Late Event Bulletin (LEB) catalog was detected and that adding our confirmed detections for the LEB catalog would more than double for regional magnitude detection. The results of this study are general, and can be applied to real-time-generated templates.

Bulletin of the Seismological Society of America, Vol. 105, No. 1, pp. 240–256, February 2015, doi:10.1785/0120140106

Initial Global Seismic Cross-Correlation Results: Implications for Empirical Signal Detectors

by D. A. Dodge and W. R. Walter

Abstract In this work, we cross-correlated waveforms in a global dataset consisting of over 310 million waveforms from nearly 3.8 million events recorded between 1970 and 2013 for two purposes: to better understand the nature of global seismicity and to evaluate correlation as a potential signal detection method. We find that about 14.5% of the events for which we have at least one waveform correlated with at least one other event at the 0.6 or higher level. Within the geographic regions where our waveform holdings are complete or nearly complete, that fraction rises to nearly 18%. Moreover, among the events for which we had one or more seismograms recorded at distances less than 12°, the fraction of correlated events was much higher, often exceeding 50%.

These results imply that global seismicity contains a large number of repeating events, that is, events that are sufficiently similar to each other to have correlated waveforms. This has important implications for signal detection, including using for unique detection in terms of automated processing. It is well known that the strength of the signal implied similarly of the sources of correlated events, to be performed ever more quickly, accurately, and lower magnitudes. If this conclusion is true, there are many implications for network operations and bulletins preparation, for example, using waveform correlation, a sparse network.

Introduction

It has long been known that seismic events produce seismic signals that are sufficiently similar to each other to have correlated waveforms. This has important implications for signal detection, including using for unique detection in terms of automated processing. It is well known that the strength of the signal implied similarly of the sources of correlated events, to be performed ever more quickly, accurately, and lower magnitudes. If this conclusion is true, there are many implications for network operations and bulletins preparation, for example, using waveform correlation, a sparse network.

Correlation detectors have also been applied with some success to earthquake aftershock sequences. Large earthquake sequences are a problem for monitoring agencies because the high rate of aftershocks makes it difficult to keep up with processing times. This is due to the sheer volume of events to be processed and the numerous false alarms produced by the automated systems underdevelopment at the time. If it is possible to use the large fraction of events to be correlated, then a seismic signal processing pipeline suitable designed to use correlations to trigger detections and perhaps other forms of ESD for regional network processing and continental global processing, because, for example, nearly all continental seismicity (99%) is within 12° of at least one International Monitoring System station.

Hasel et al. (2013) Here, we test correlation as a type of empirical signal detector (ESD), a term coined by Juska et al. (2013) to refer collectively to pattern matching detectors such as correlation and subspace detectors (e.g., Harris, 2006; and most references therein).

Correlation detectors have also been applied with some success to earthquake aftershock sequences. Large earthquake sequences are a problem for monitoring agencies because the high rate of aftershocks makes it difficult to keep up with processing times. This is due to the sheer volume of events to be processed and the numerous false alarms produced by the automated systems underdevelopment at the time. If it is possible to use the large fraction of events to be correlated, then a seismic signal processing pipeline suitable designed to use correlations to trigger detections and perhaps other forms of ESD for regional network processing and continental global processing, because, for example, nearly all continental seismicity (99%) is within 12° of at least one International Monitoring System station.

Harris and Dodge (2011) have used correlation in combination with subspace detectors in an automated system to track events in an aftershock sequence. They demonstrated a



Motivation (2)

- **However, the basic calculation (waveform correlation) is computationally expensive, so either the size of waveform archive used has been modest...**
 - ~3 templates, NK Test Site Monitoring (Schaff et al., 2012; Zhang & Wen, 2014)
 - ~200 templates, small aftershock sequence (Dodge & Harris, 2011)
- **...or complex distributed computing systems have been used**
 - ~5000 templates, broad regional monitoring (Slinkard et al., 2014 – used SNL developed distributed)
 - ~500M templates**, global scale research (Dodge & Walter, 2015)

***this study did NOT perform scanning of continuous data but rather performed cross-correlations of windowed signal detection waveforms in the LLNL GNEMRD database*



Objective

- Utilize Approximate Nearest Neighbor (ANN) methods developed in other domains (e.g. searching for images) to enable searching of very large waveform archives without requiring a complex distributed computing system.



Test Case

- **Our test case: 14 Years of IDC LEB MKAR Signal Detections (2002-2014)**
 - Archive: 2002-2013
 - 308K analyst reviewed arrivals (240K events)
 - Test Set: 2014
 - 26K analyst reviewed arrivals
 - 60% teleseismic P, 14% regional
 - Used only MK01/SHZ channel (no array processing)
 - 0.5-5.0 Hz bandpass filter (generic multi-purpose)



Results

- **20x speedup** finding highly correlated waveforms
- How? In exchange for finding 80% of the events in the archive
- Quick to set up: ANN Index for 308K waveforms builds in 180 seconds
- Quick to query: ~30msec for each search, ~1sec to calculate correlation scores
- All on a standard desktop computer



Approximate Nearest Neighbor Search Methodology



Approximate Nearest Neighbor (ANN) Search

Return an approximation of the true nearest neighbor set from an archive of data ...

In exchange for large increases in query speed over standard linear search

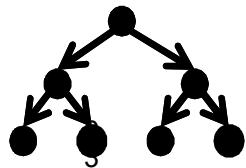


ANN Application

- Best suited for tasks where the database is expected to have large number of useful neighbors and the specific subset of those returned is not critical
 - Quickly eliminating events from common seismic zones or known mining operations
- Not suited for “Needle in a Haystack” Problems
 - Identifying NK Test Events

ANN Index Types

Tree Based



Euclidean Similarity Only

Hashing Based



$h(x)$

010100101010010



Hash Table

ANN Index Types

Challenge for Seismic Waveforms

Tree Based

Hashing Based

Find projection function $p(x)$

$euclidean(p(x), p(y)) \propto correlation(x, y)$

Develop hash function

$h(x)$

KPCA Data Projection



Raw Data (19s at 40hz) = 760 dimensions



Kernalized Principal Component Analysis

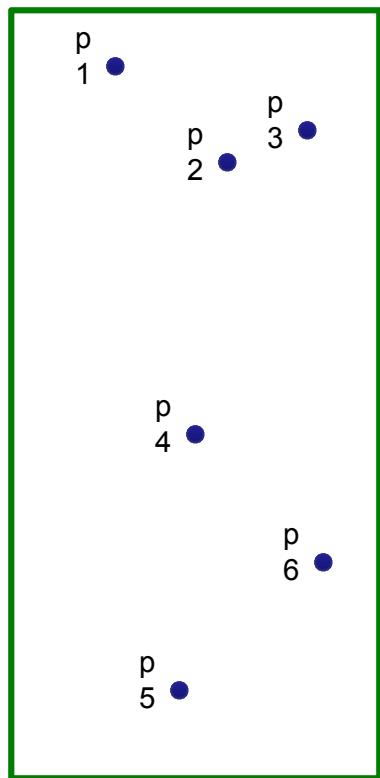
$euclidean(kpca(x), kpca(y)) \propto correlation(x, y)$
dimensionality reduction

Projected Data = 200 dimensions

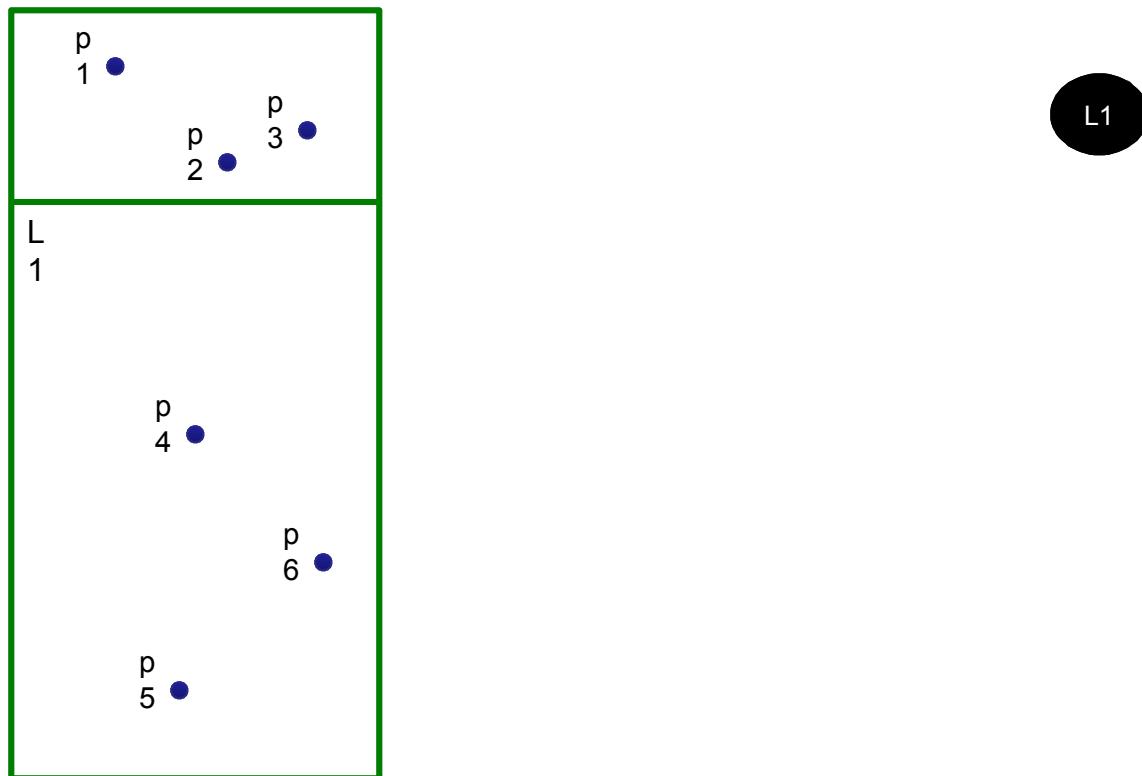
KD-Trees

- K-Dimensional Tree (KD-Tree)
 - Partition a large database of N k -dimensional objects using a tree structure
 - $\log(N)$ search time

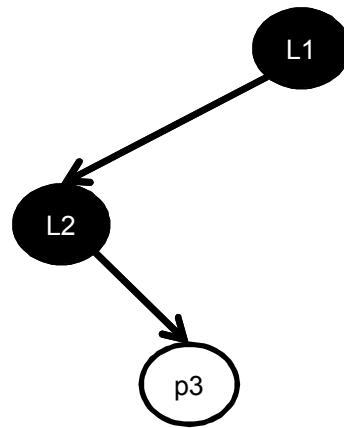
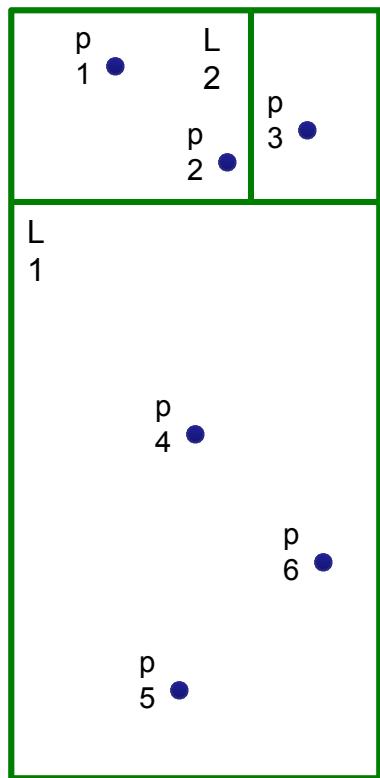
Building a KD-Tree



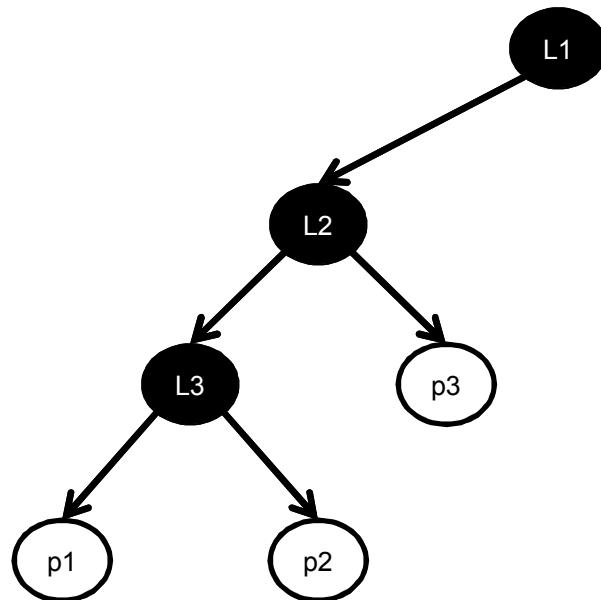
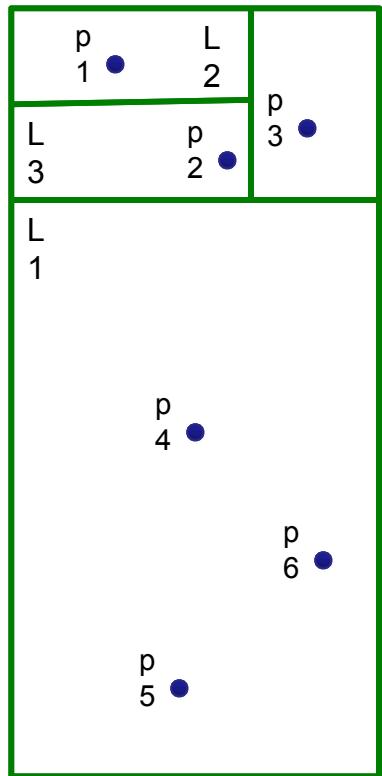
Building a KD-Tree



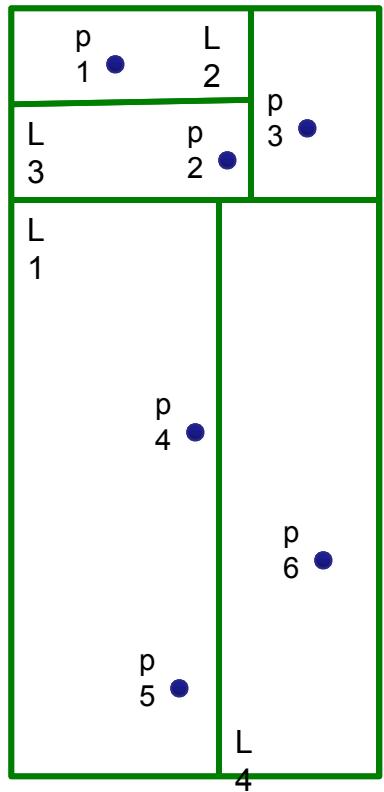
Building a KD-Tree



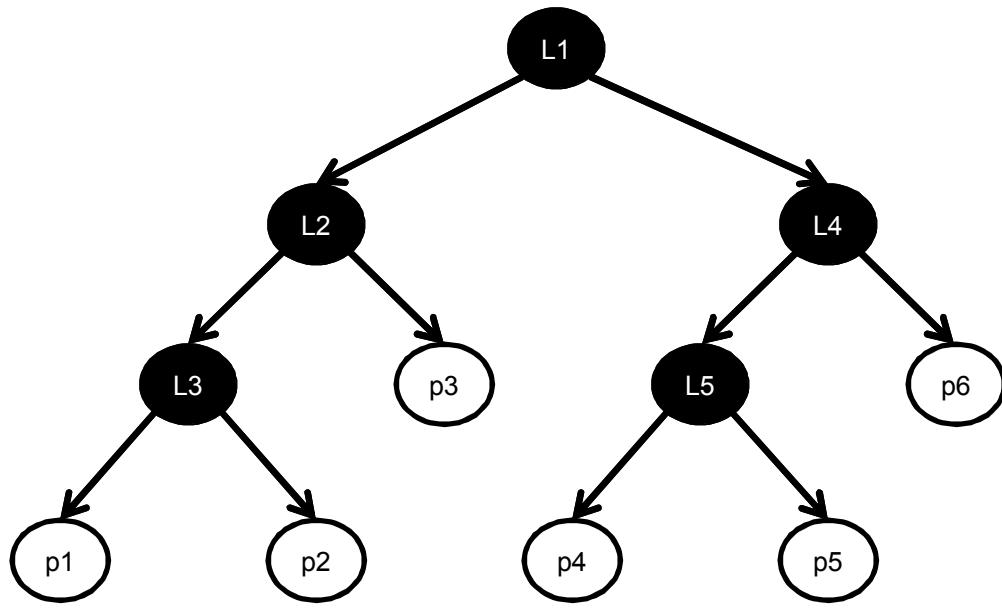
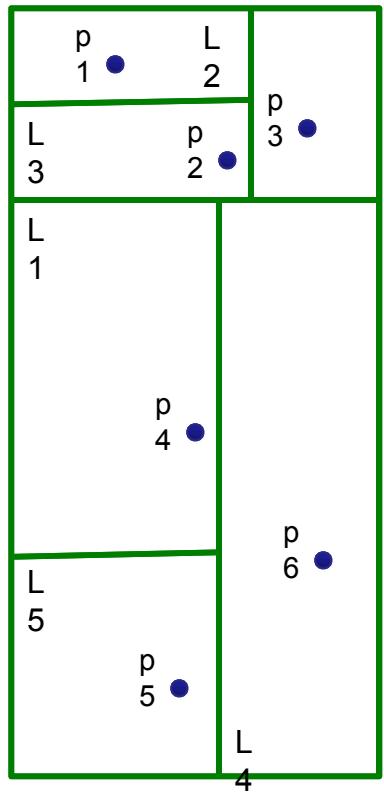
Building a KD-Tree



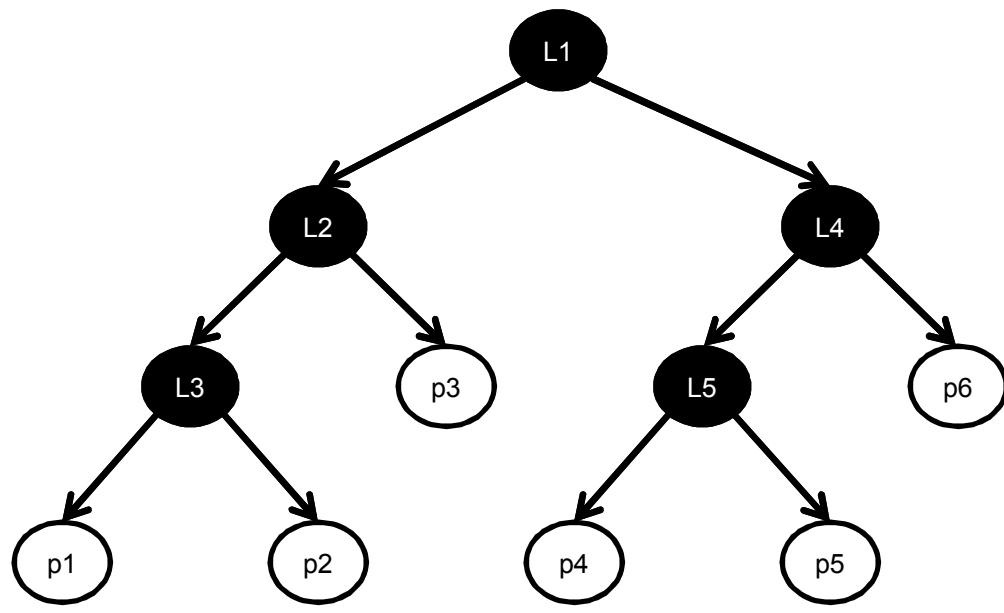
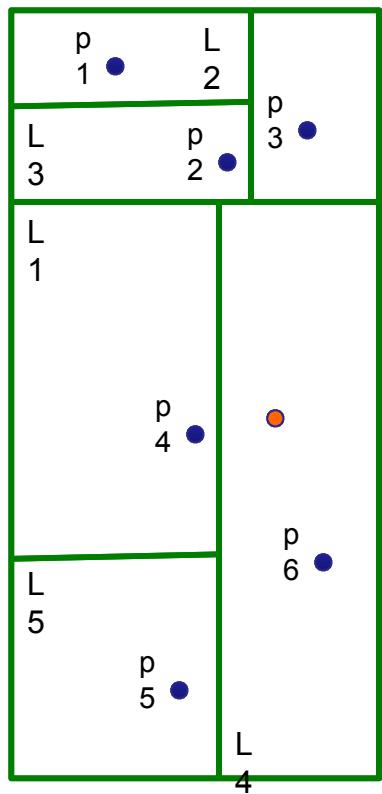
Building a KD-Tree



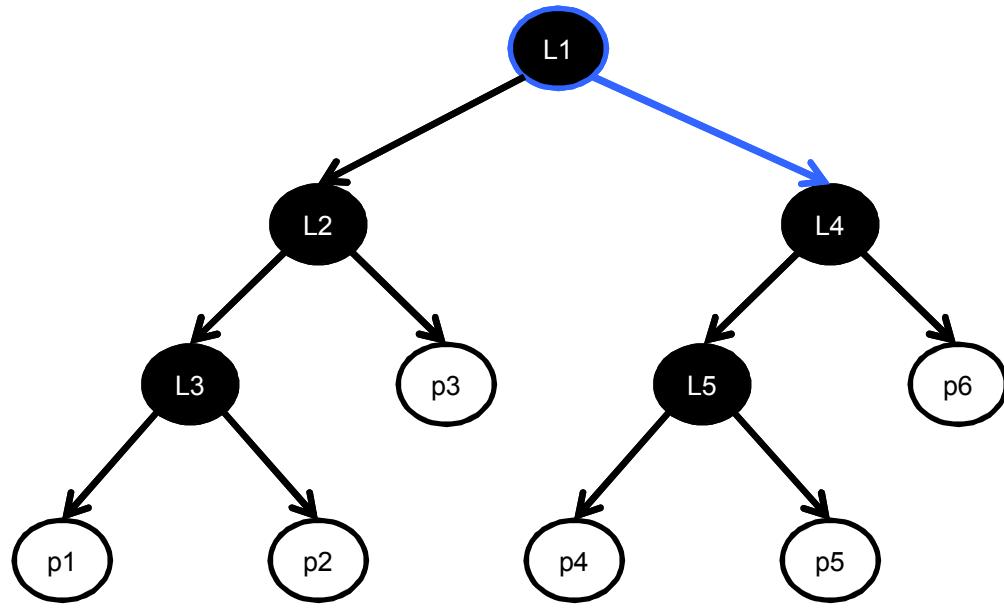
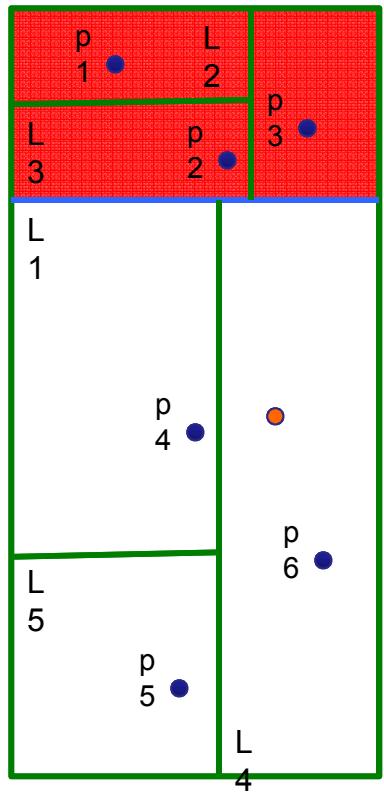
Building a KD-Tree



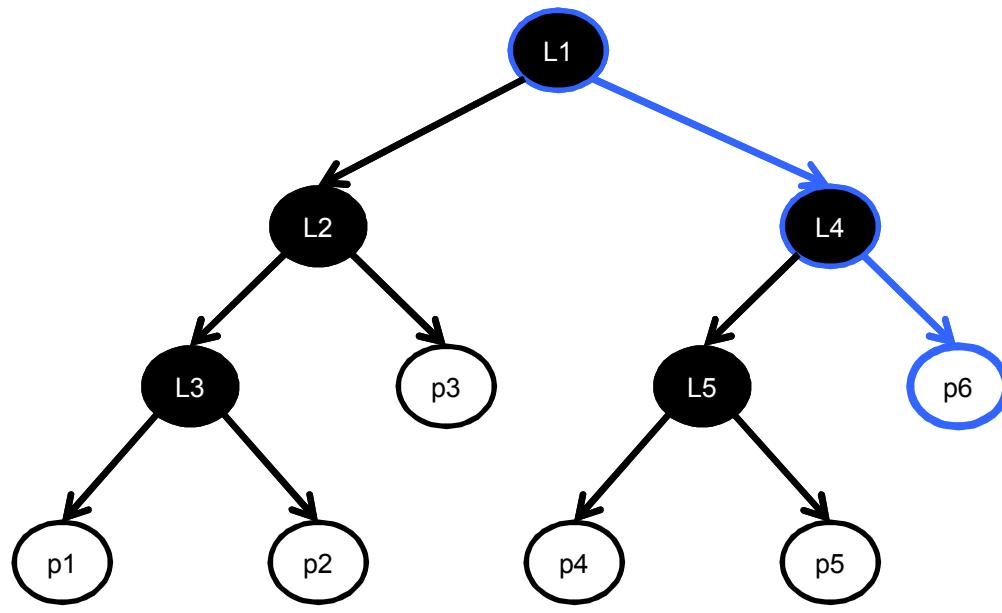
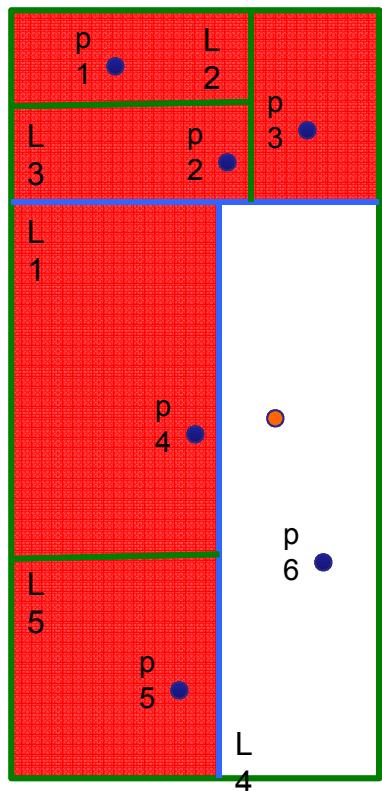
Querying a KD-Tree



Querying a KD-Tree

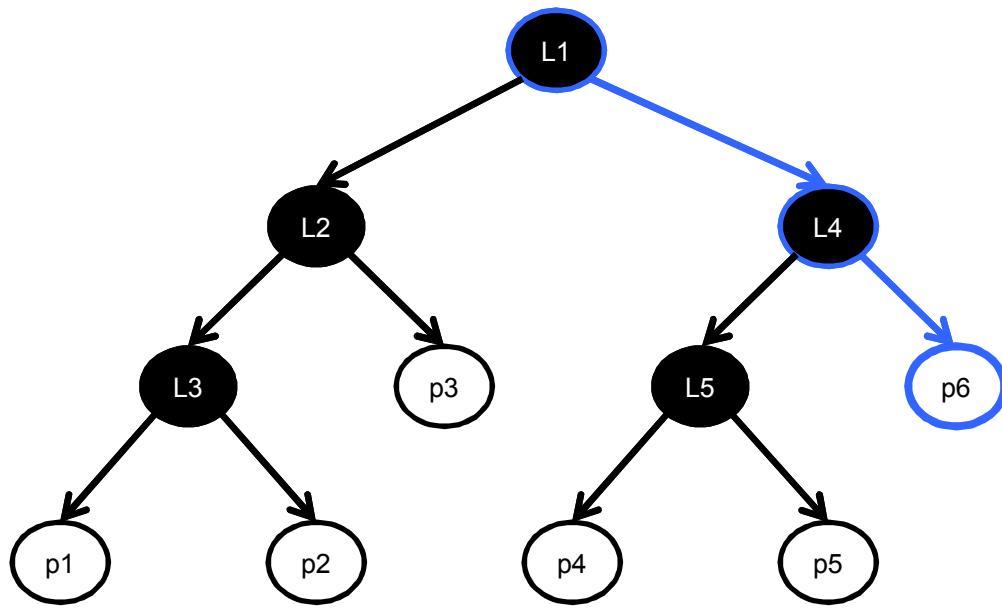
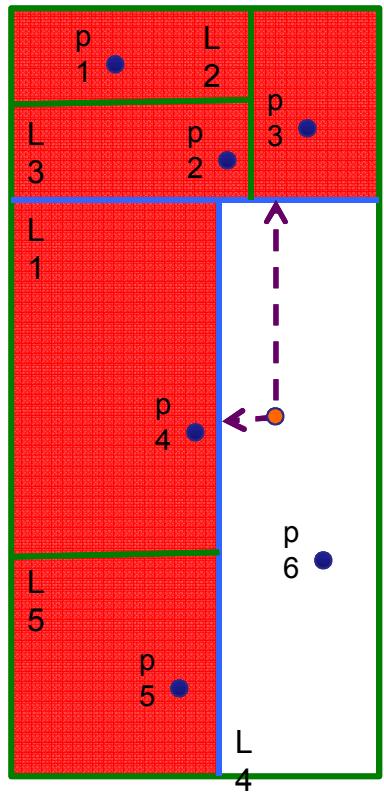


Querying a KD-Tree



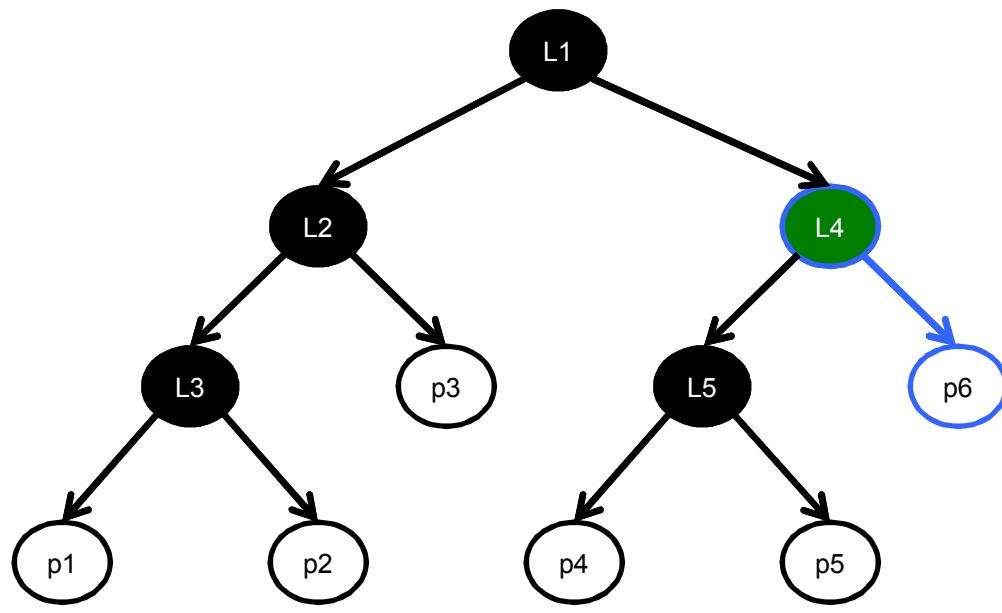
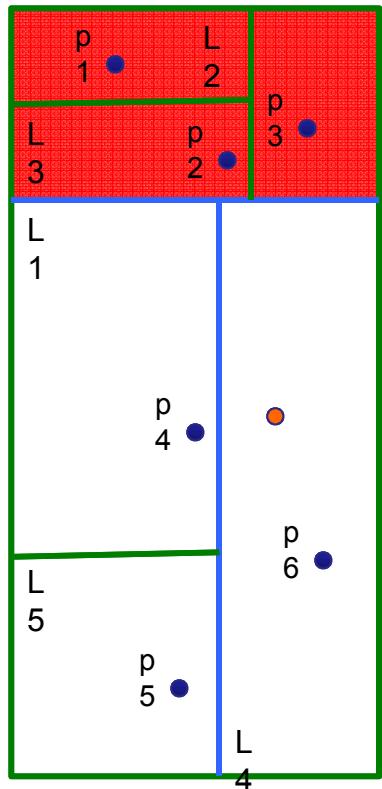
Found Neighbors: p6

Querying a KD-Tree: Best Bin First



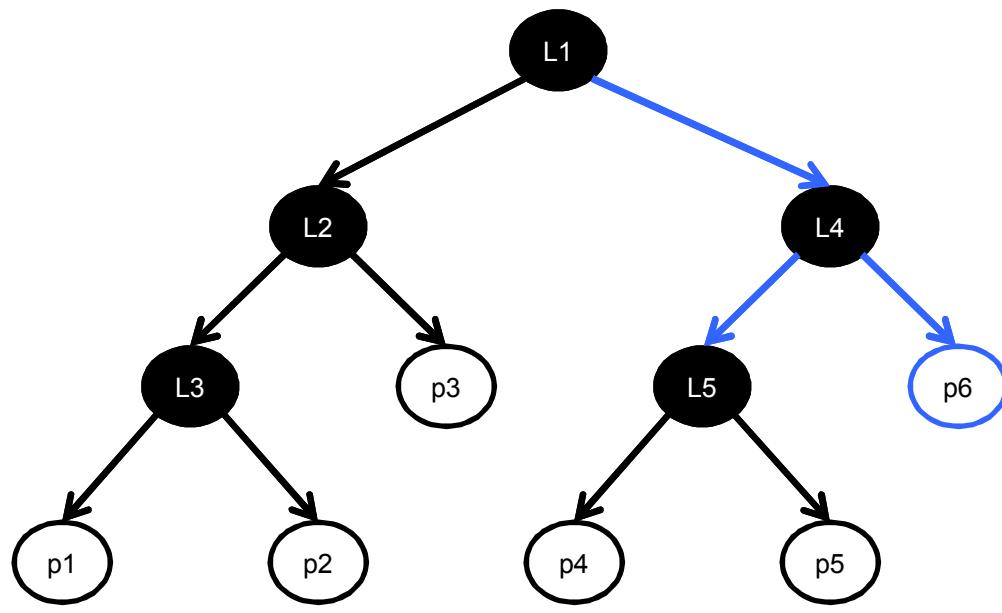
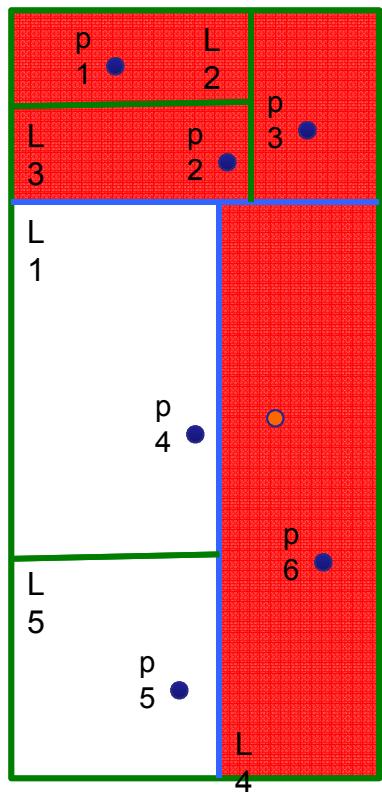
Found Neighbors: p6

Querying a KD-Tree: Best Bin First



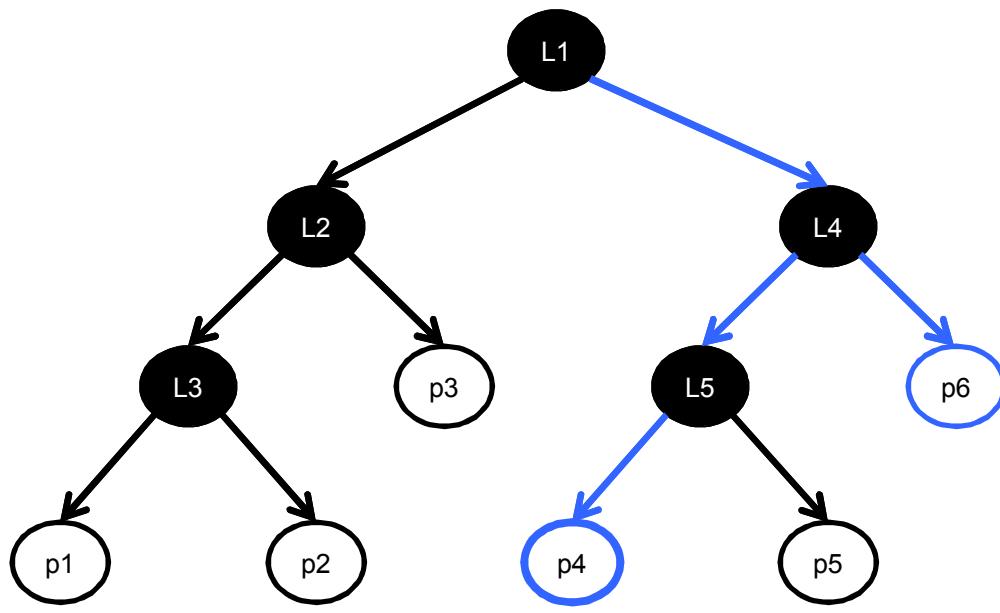
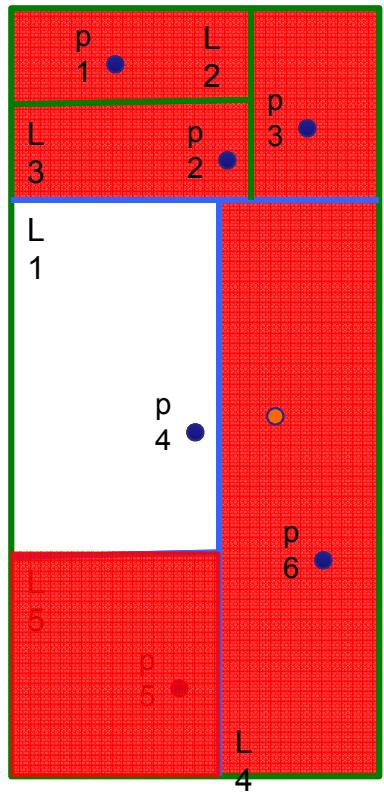
Found Neighbors: p6

Querying a KD-Tree



Found Neighbors: p6

Querying a KD-Tree



Found Neighbors: p6, p4



KD-Tree Forests

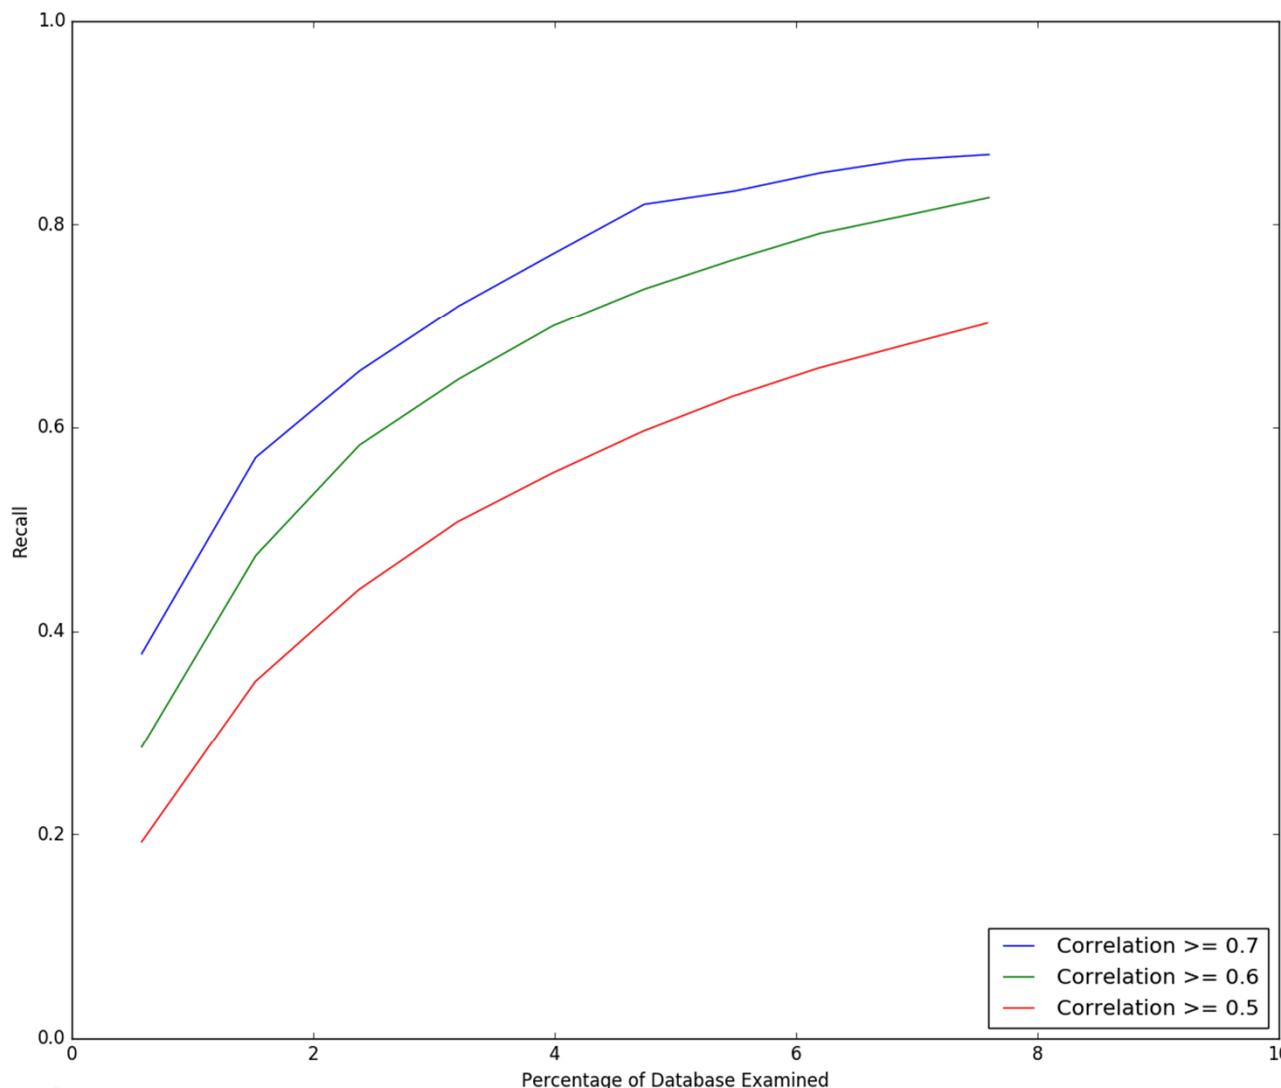
- KD-Tree Forest: build an index from a group of randomly varied KD-Trees and query them collectively
 - During each split decision, randomly select from the top n highest variance dimensions
 - During query, perform Best-Bin-First collectively across all trees
- For seismic case, we use forests with 100 trees



ANN Queries in Practice

- Due to the approximate nature of ANN, a query into the index must:
 - **Over-sample:** Return a set of candidate nearest neighbors larger than what is need to perform the given task
 - **Prune:** Compute the true correlation between the query and all members of the candidate set, sort, and discard those that do not meet some quality threshold

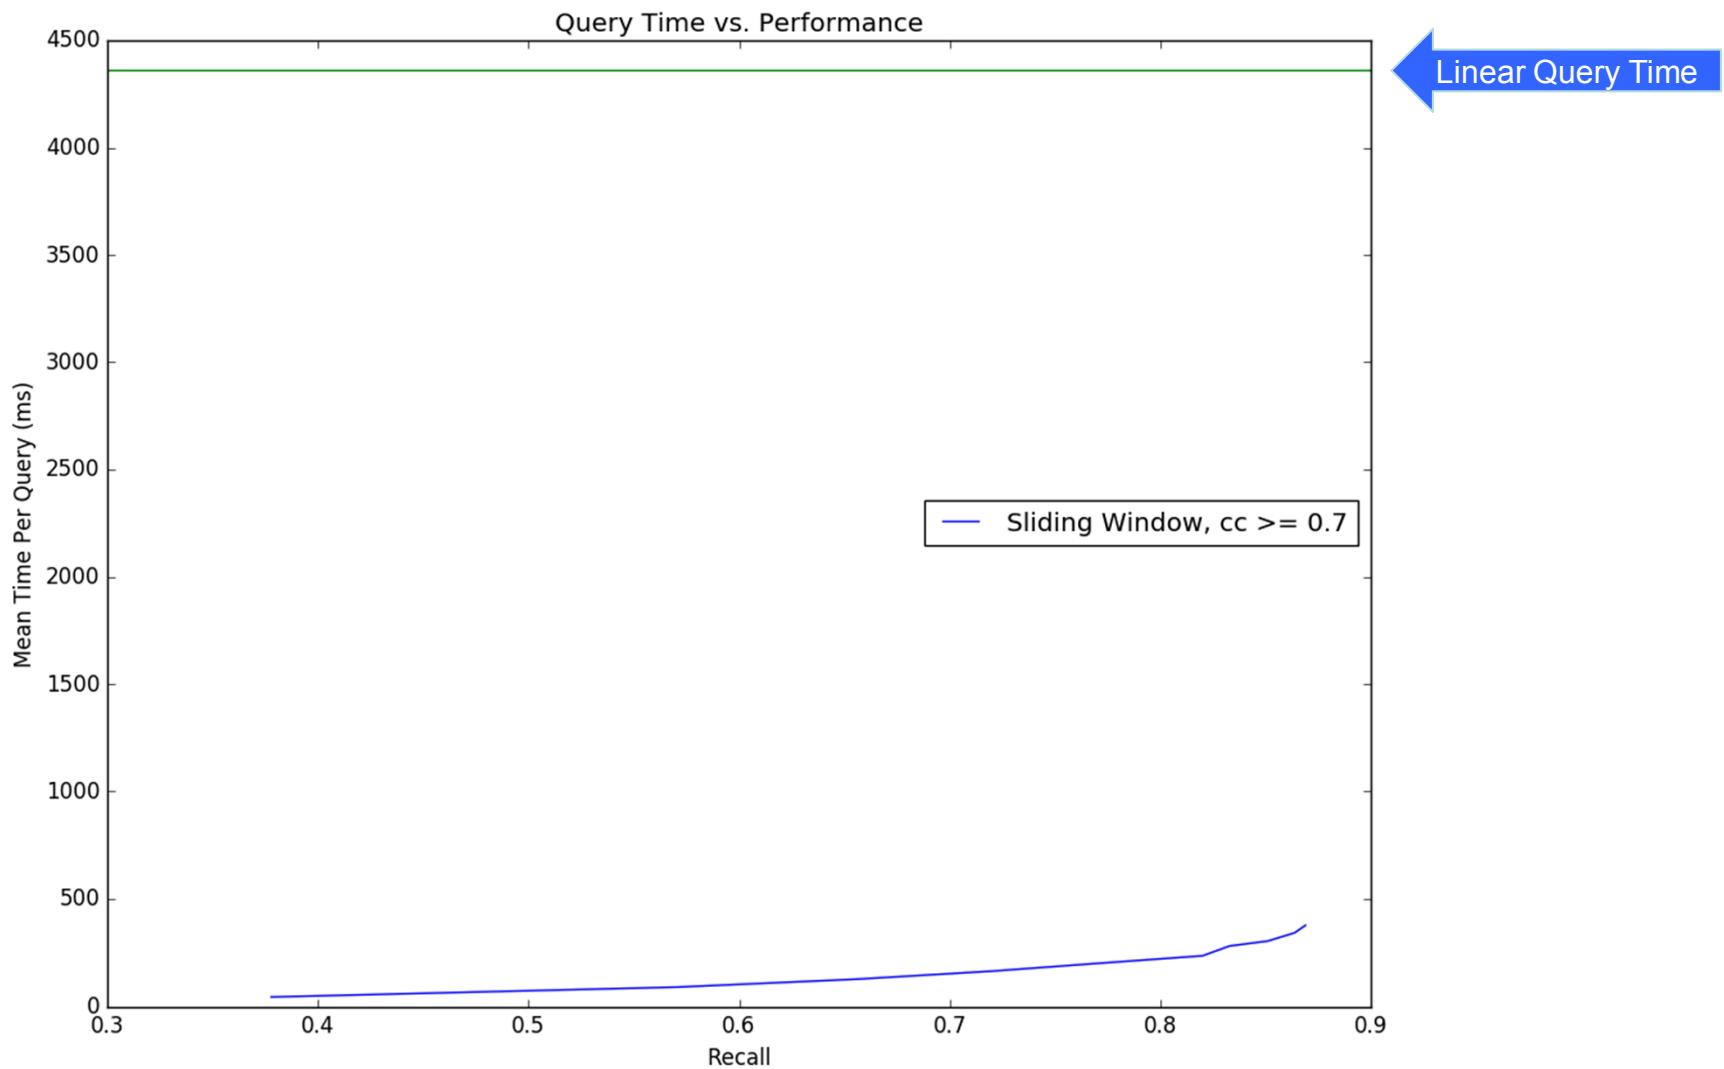
ANN Search: Performance Evaluation



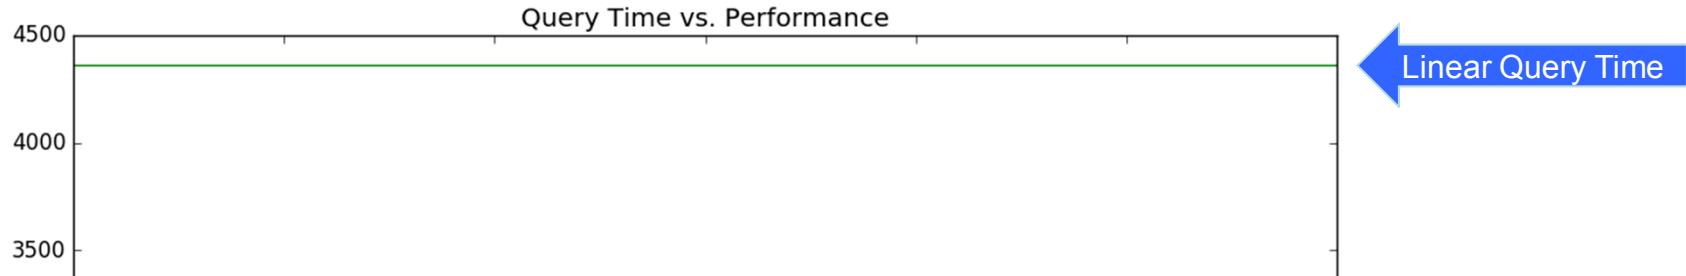
Performance metric is *recall*: the percentage of high quality neighbors returned by ANN query



ANN Search: Performance Evaluation

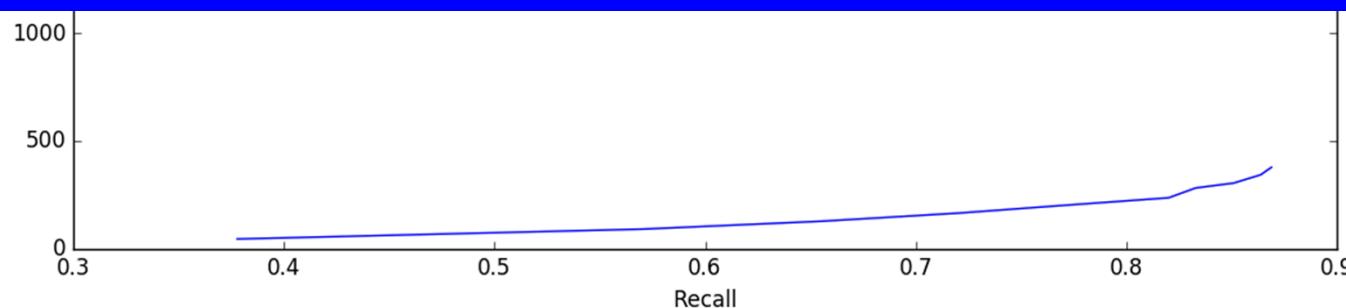


ANN Search: Performance Evaluation



Pruning is not always required as some tasks can be completed by only looking at the metadata of the returned NN candidates.

Without the pruning step, ANN waveform queries are 115x faster than linear ones

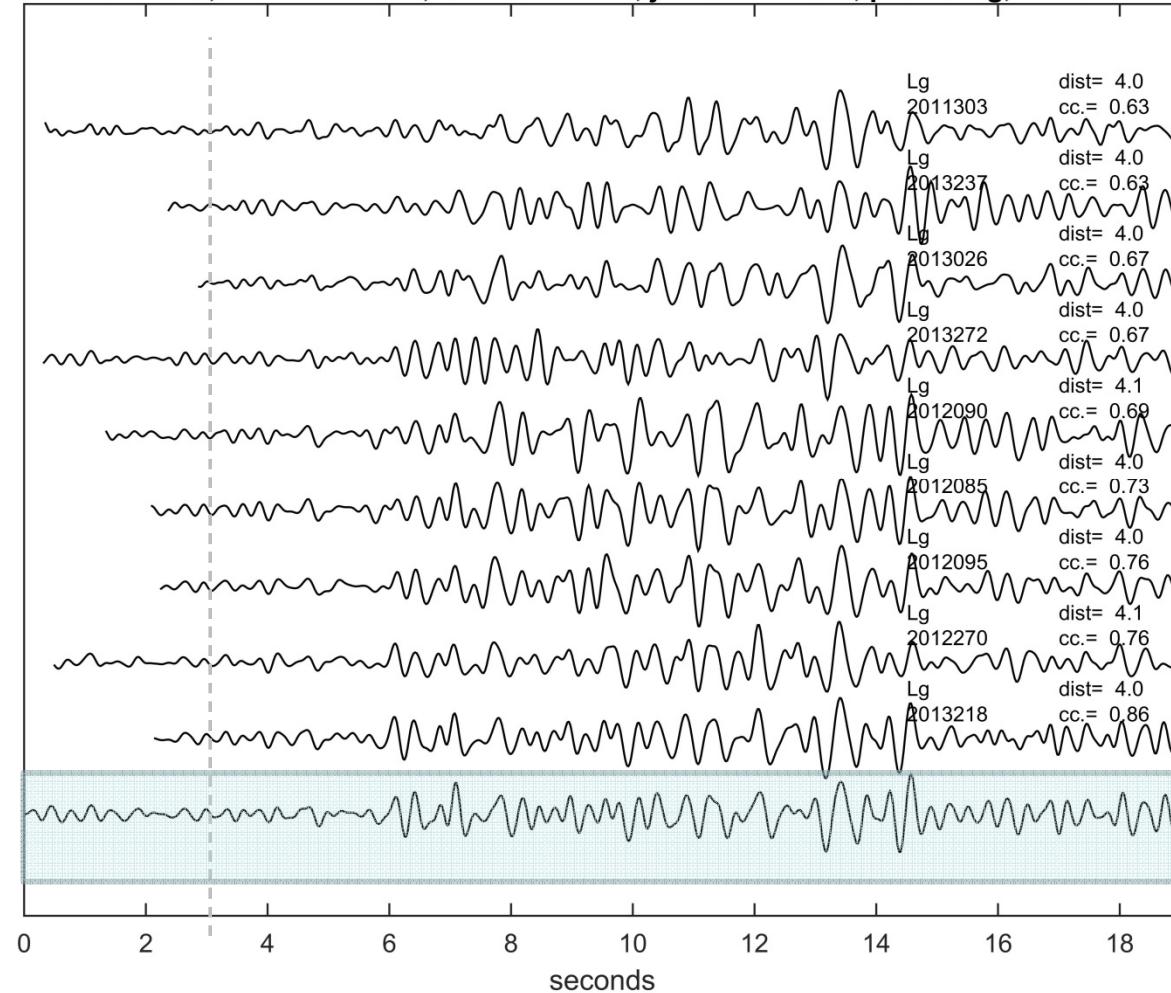




Test Results

Regional Example #1 (Lg)

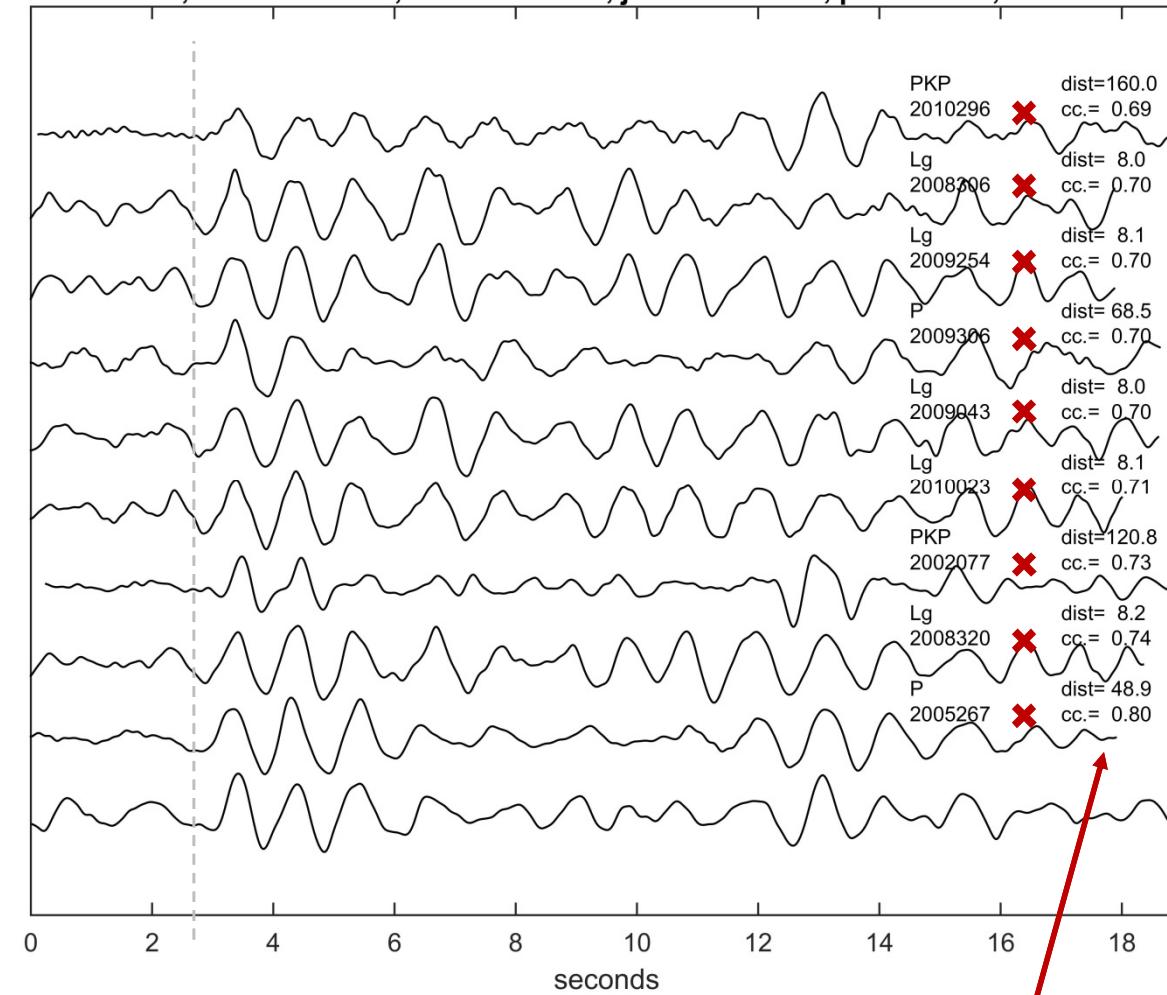
wfmid=539, orid=10440936, arid=92021278, jdate= 2014008, phase=Lg, delta=4.0



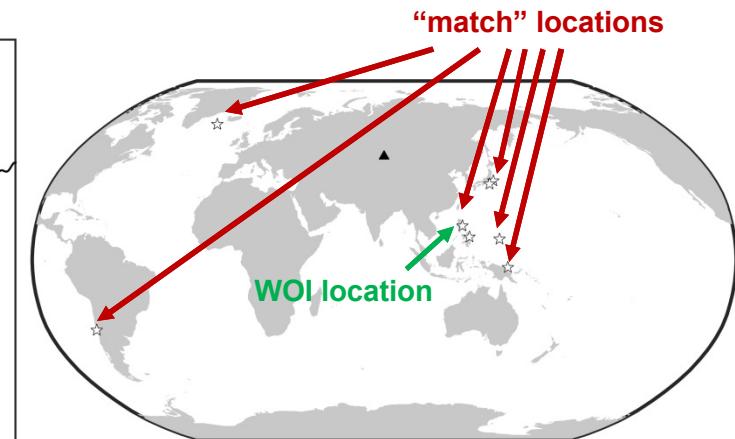
- Waveform to be identified is shown at bottom
- Matches shown above, in decreasing order of correlation
- Match pool span 2011-2013
- Mining explosions?

Screening Example #1: Rejected “Match”

wfmid=3, orid=10414339, arid=91845654, jdate= 2014001, phase=PcP, delta=41.6



Very high correlation score!



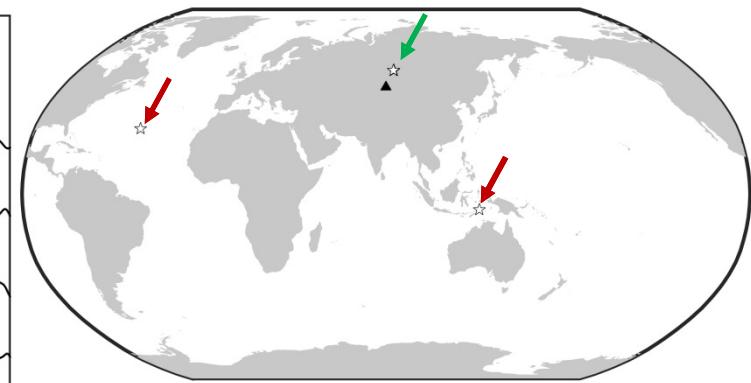
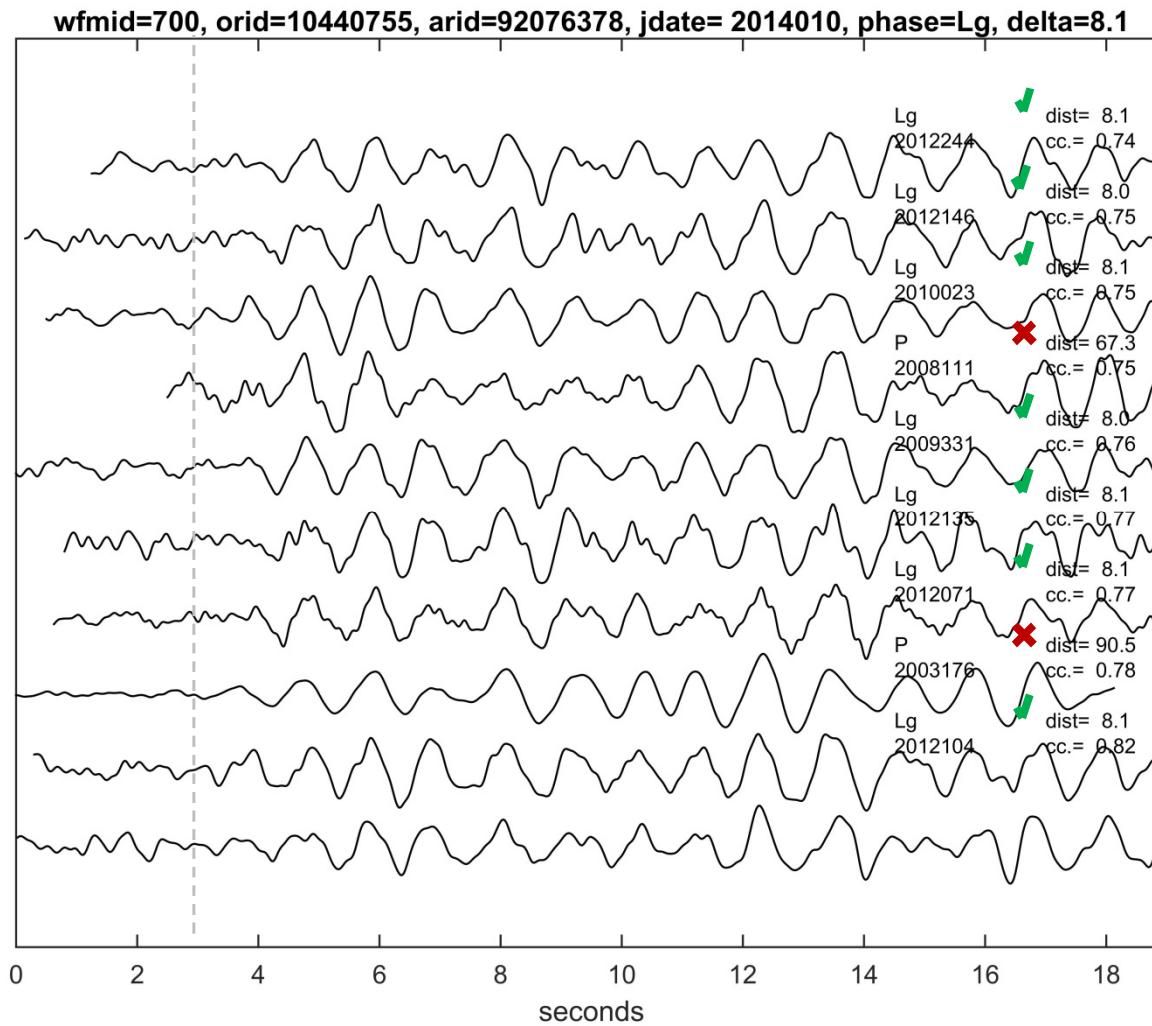
- Only showing top 9 (there are 136 with $cc > 0.60$ in this case)
- None of these is in the right location, but we wouldn't know that
- No consensus in location among match pool, so we reject the match



Match Results Metadata Screening: An Essential Addition

- High correlation score isn't enough
- Can use *consistency* in event location of match pool to decide whether to accept match
- Developed a single algorithm based on event location consistency of 3 of the top 4 results
- More sophisticated approaches could be tried, but this worked quite well for this first simple study.

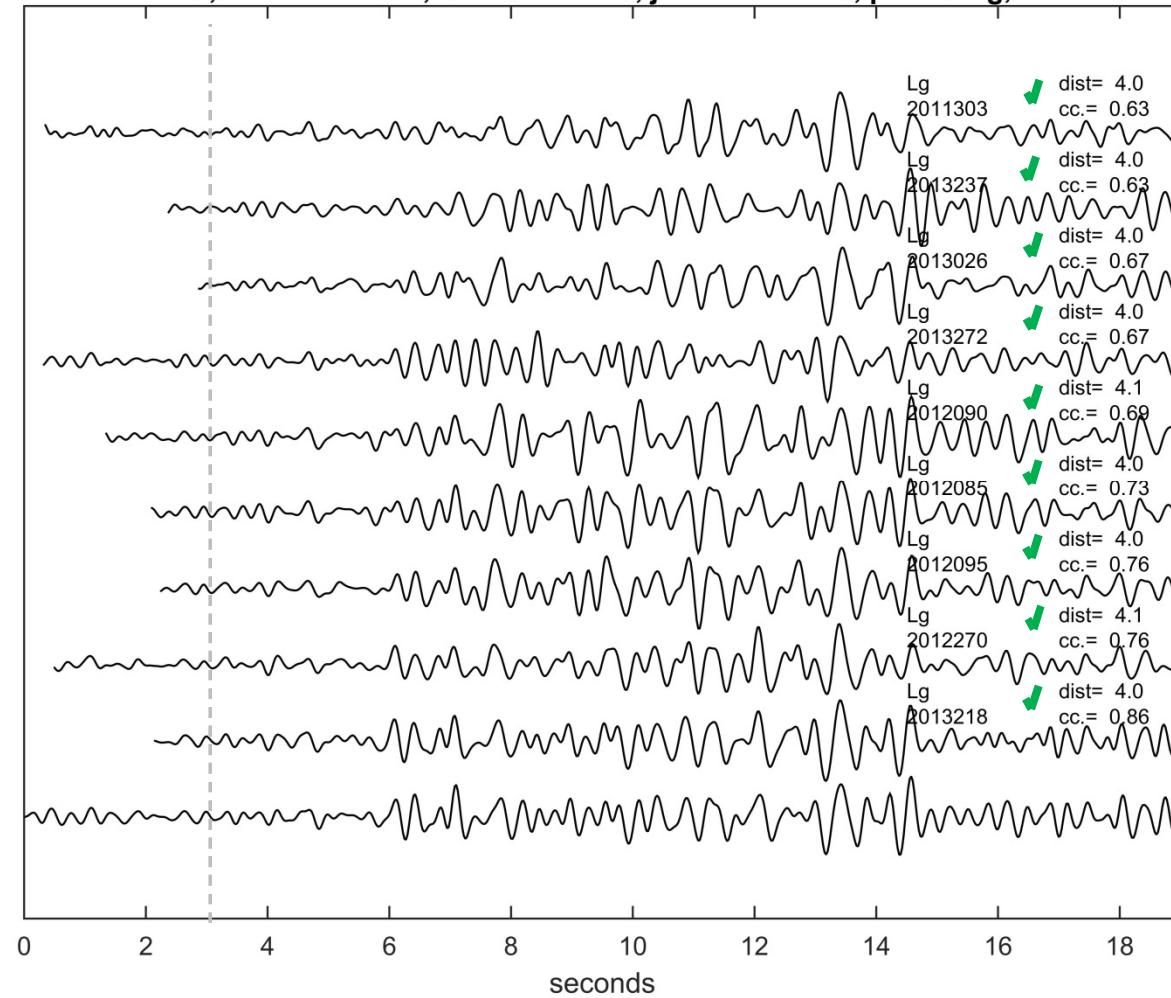
Screening Example #2: Accepted Match”



- Results include correct and incorrect matches
- This is a confirmed match: algorithm tolerates some dissent (3 of top 4 have excellent agreement)

Regional Accepted Match Example #1 (Lg)

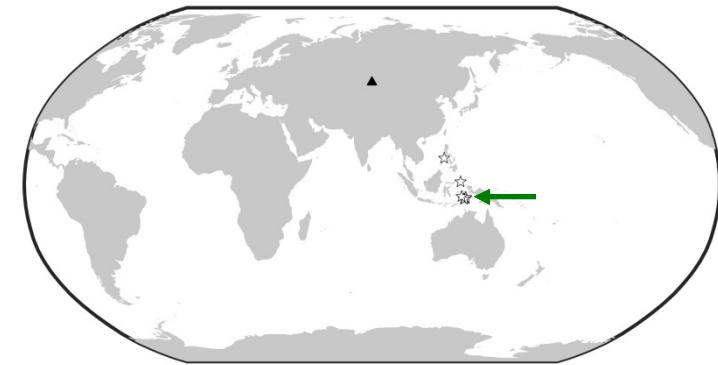
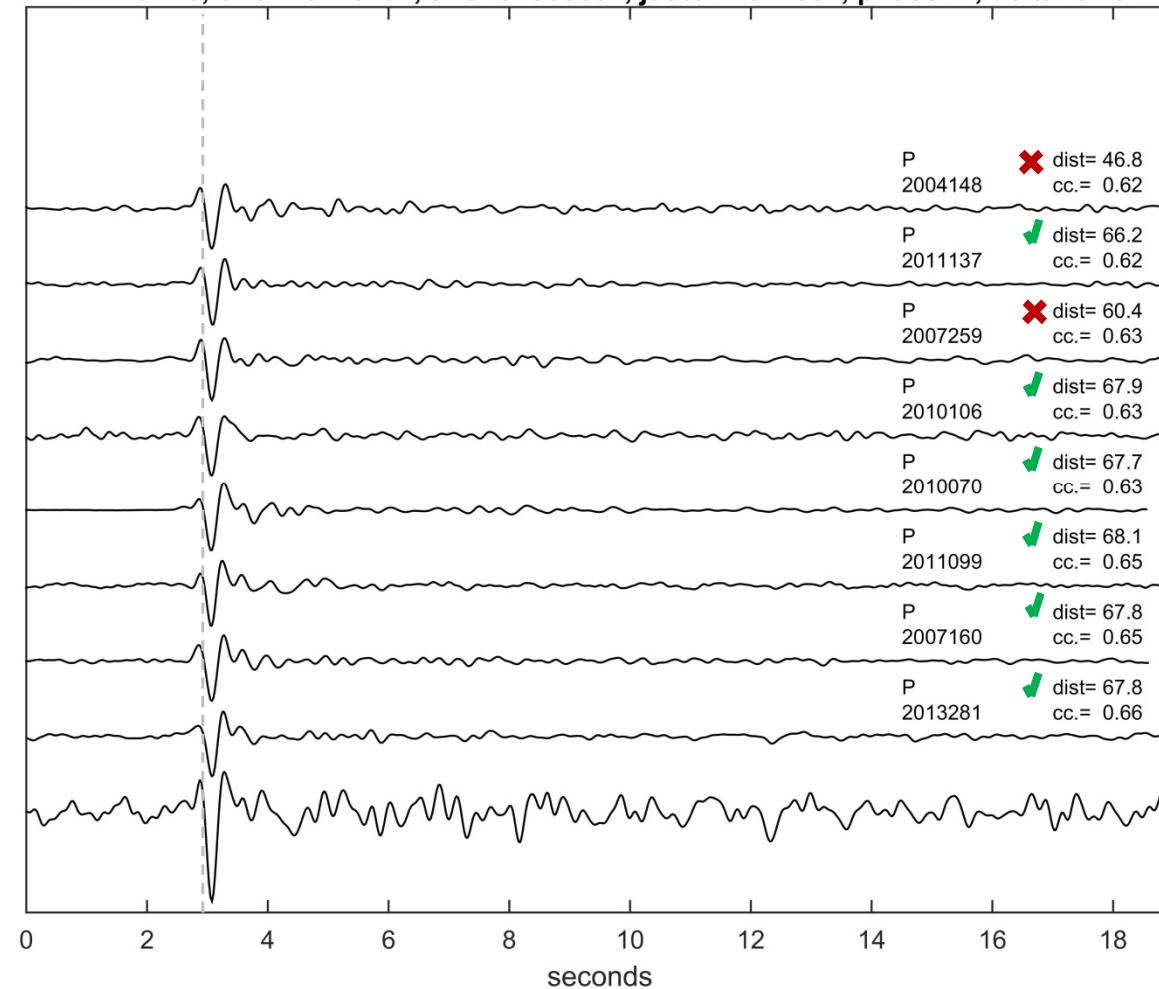
wfmid=539, orid=10440936, arid=92021278, jdate= 2014008, phase=Lg, delta=4.0



- Waveform to be identified is shown at bottom
- Matches shown above, in decreasing order of correlation
- Match pool span 2011-2013
- Mining explosions?

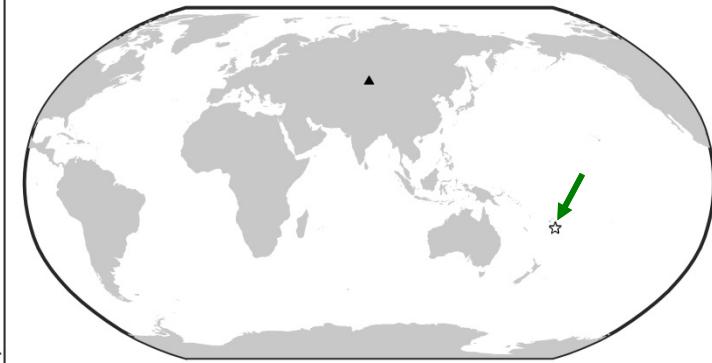
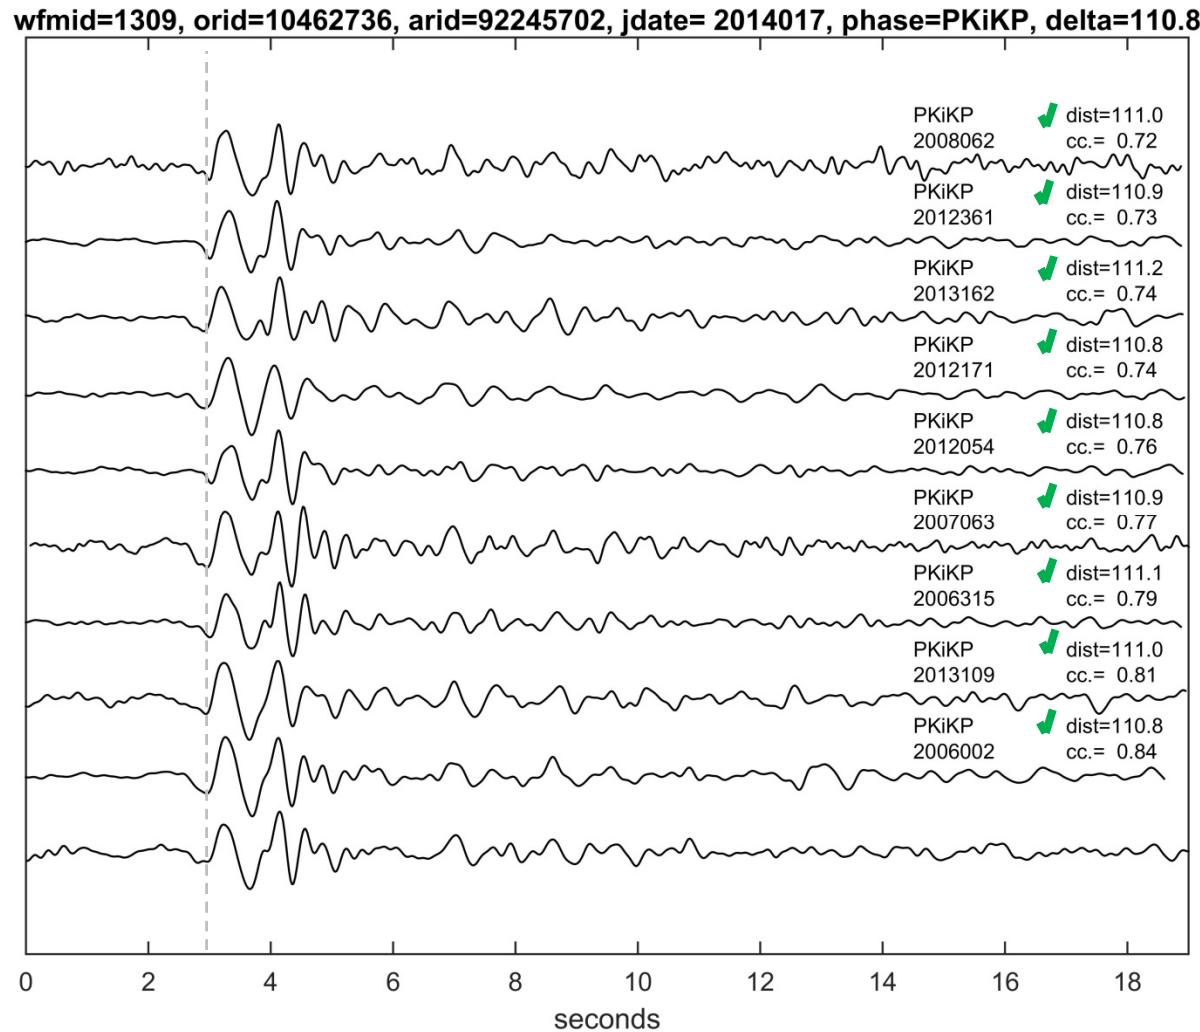
Teleseismic P Accepted Match Example #1

wfmid=76, orid=10425131, arid=91866507, jdate= 2014001, phase=P, delta=67.9



- Much shorter duration, simpler signal, but still matched!
- Pool of good matches much smaller (only 8): earthquakes have longer recurrence interval than mine explosions
- Match pool spans 2004-2013

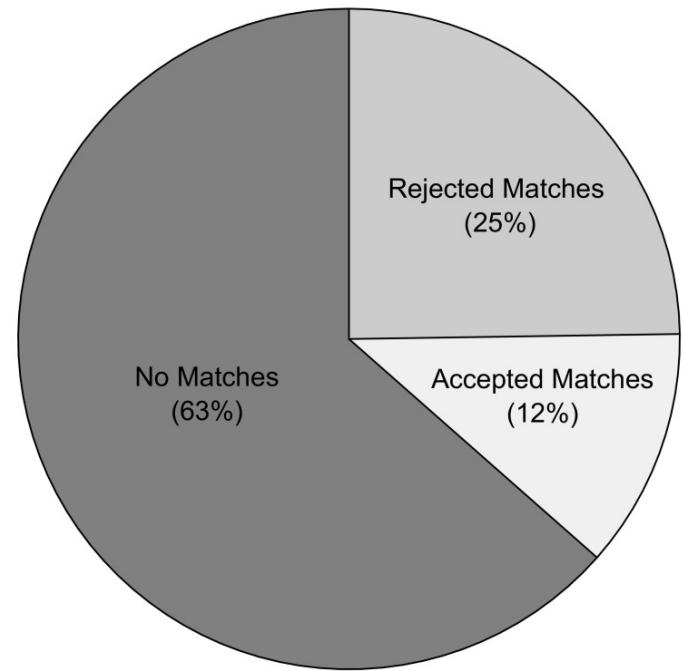
Teleseismic PKiKP Accepted Match Example #1



- Very complex signal is unmistakable (high confidence match)
- Match pool spans 2006-2013

Match Results Summary (1)

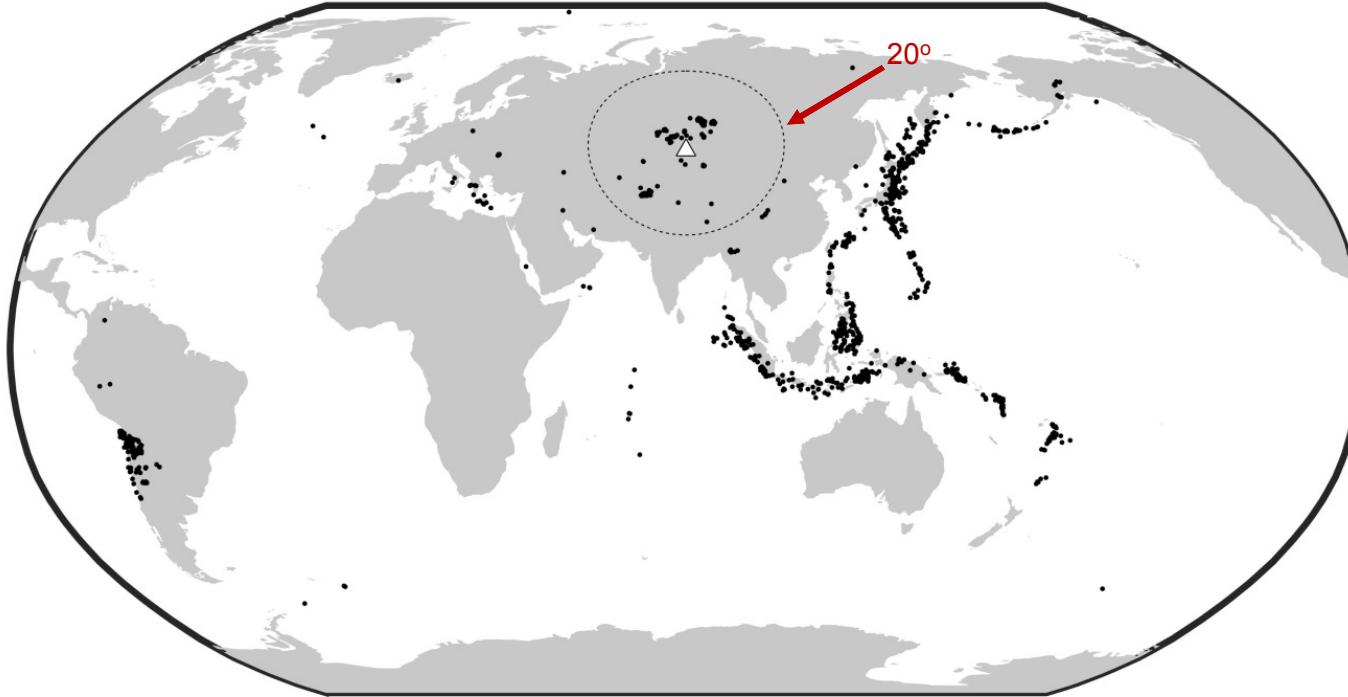
- 63% of 2014 signal detections could not be identified (no matches or 1 match)
- Of remainder, 25% were rejected due to lack of event location consistency
- Analysis of 12% accepted matches:
 - 68% of those validated using LEB info
 - 93% of regional matches (Pn, Pg, Sn, Lg) validated
 - 67% of teleseismic first P validated, (important because 60% of all LEB associated detections are tele first P)
 - 71% of core phases validated (PKPbc, PKB, PKiKP))
 - Poorer performance for secondary phases (42% PcP matches validated)



2014 signal identification



Match Results Summary (2)



- Majority of matches are at teleseismic distances because majority of events are at teleseismic distances
- Nearly all teleseismic matches are likely to be earthquakes
- Having a long time baseline for template archive is very important for matching earthquakes (much slower recurrence rate than mining events)





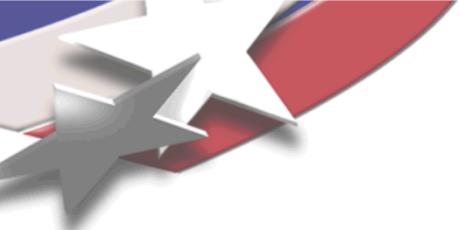
Summary

- **Motivation** for work was observation that computation cost of waveform correlation was limiting size of signal archives that could be searched
- **Objective** was to see if ANN methods could facilitate searching of very large archives
- **Findings**
 - Our KPCA + KD Trees approach allows us to search archives at least as large as the total number of signals recorded by one of the most active stations in the IMS network (~300K, algorithm can easily handle much larger)
 - No specialized distributed computing system needed: uses single COTS PC (albeit a fairly powerful one)
 - Adding a match pool screening step based on consistency of metadata was crucial for getting good results
 - Overall the proportion of detected signals that could be matched was modest (12% total, 8.6% validated), but impact on automatic processing could still be very significant
 - Long archive time baselines very important for matching earthquakes



Future work

- **Frequency dependence:** Build indexes for several overlapping narrower bands to try to better enhance the signals from particular event locations.
- **Array processing:** Build indexes for multiple channels and combine the results to take advantage of the coherency and move-out of signals across the array. This should dramatically reduce the number of false event locations in the ANN set of potential matches.
- **Signal Detector:** Apply directly as a signal detector rather than as an identifier of signals detected by another algorithm. Already the algorithm performs better than real-time, so could be used as an operational detector. We are currently working on implementing it in a distributed computing framework to make it fast enough to perform repeated testing on long time intervals (e.g. several months).



Backup slides



Archive Projection: Reducing the Size of the Kernel Matrix

- When the database is really large, we save time by building the kernel representation from r randomly sampled representatives.

$r \ll n$ (e.g. 500 out of 300,00)

$$K = \begin{bmatrix} k(x_1, x_1), k(x_1, x_2), \dots, k(x_1, x_r) \\ k(x_2, x_1), k(x_2, x_2), \dots, k(x_2, x_r) \\ \dots \\ k(x_r, x_1), k(x_r, x_2), \dots, k(x_r, x_r) \end{bmatrix} \quad r \times r \text{ matrix}$$



Archive Projection: Transforming the Kernel Matrix

- For next step, K must be symmetric, positive semi-definite
- $k=corr(x,y)$ does not meet this requirement, but we can transform the correlation coefficients to fix this problem:

$$k = e^{(1 - corr(x, y))}$$



Archive Projection: PCA in Kernel Space

- KPCA: Applying Principal Component Analysis to the Kernel matrix (K) creates a matrix (W) that can be used to project the archive to a form such that correlation scores have been mapped to Euclidean distance, i.e. correlated items can be found via searching for Euclidean nearest neighbors

ANN Search: KD-Tree Forests

- **KD-Tree Forest:**

Build an ANN index from a group of randomly varied KD-Trees and query them collectively

Advantages	Disadvantages
Log(d) query time, easy to implement, memory efficient	Only works for Euclidean similarity

Our interest is Waveform Correlation

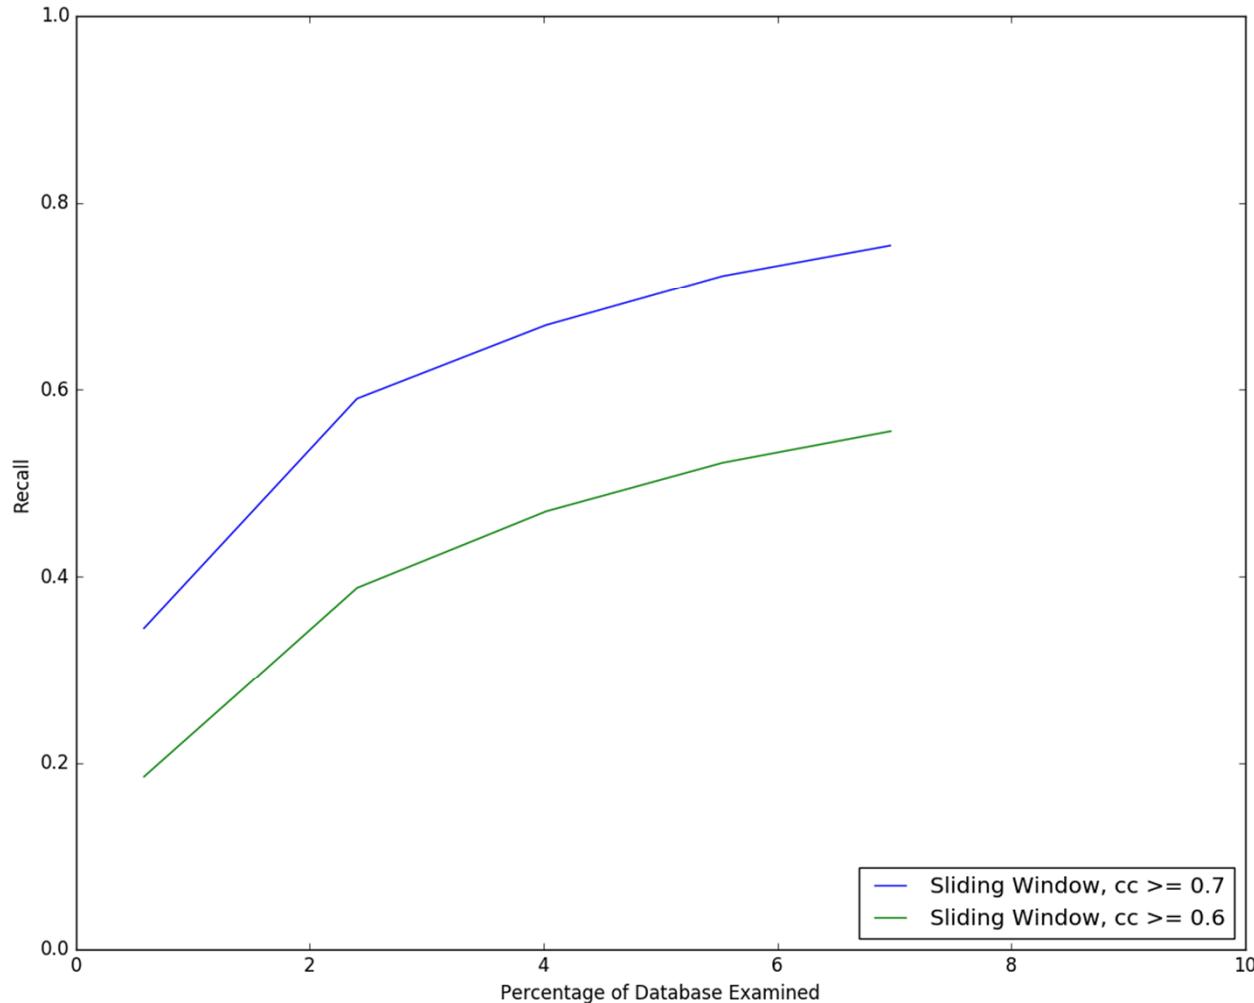


Archive Search: KD-Tree Forests

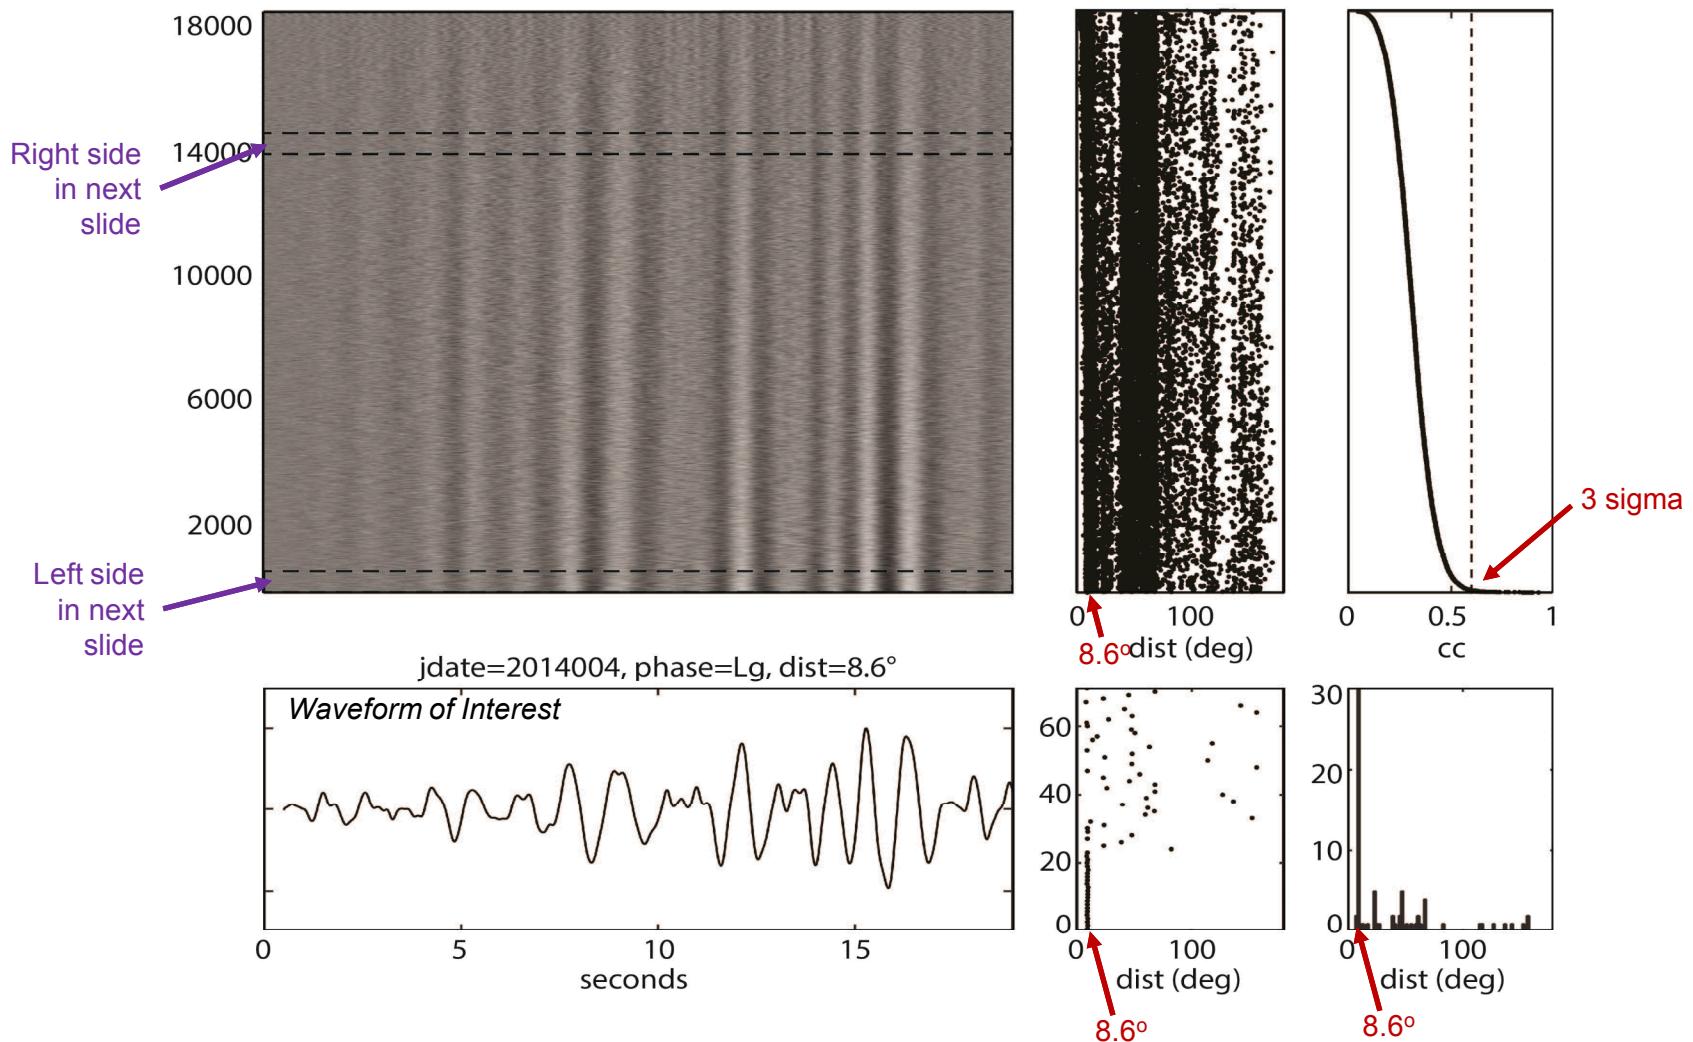
- How do we create a set of randomly varied trees?
 - During each split decision, randomly select from the top n highest variance dimensions
- How do we search the forest?
 - Search all trees until a leaf is encountered
 - During initial search, keep a sorted list of decision points based on how close the query was to the mean
 - Continue search by walking the list and taking the opposite path down to a leaf (re-visit most marginal decisions)

ANN Search: Performance Tests

Sliding window performance based on quality of neighbors

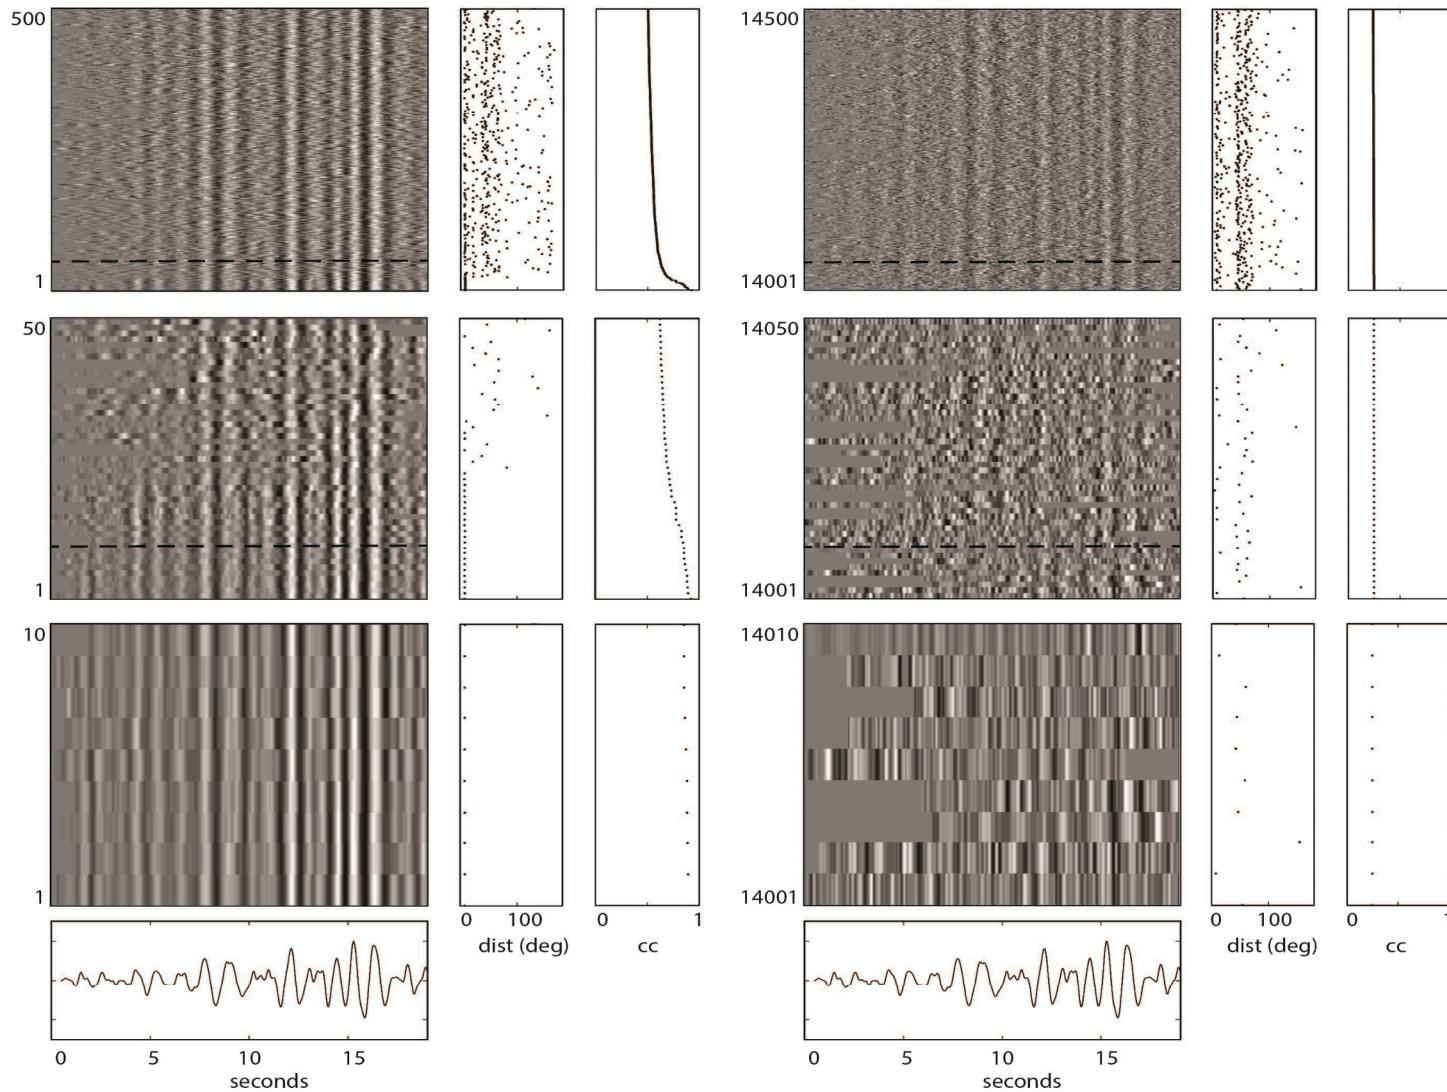


Example of KLSH results for a single waveform (Lg, 8.6°)



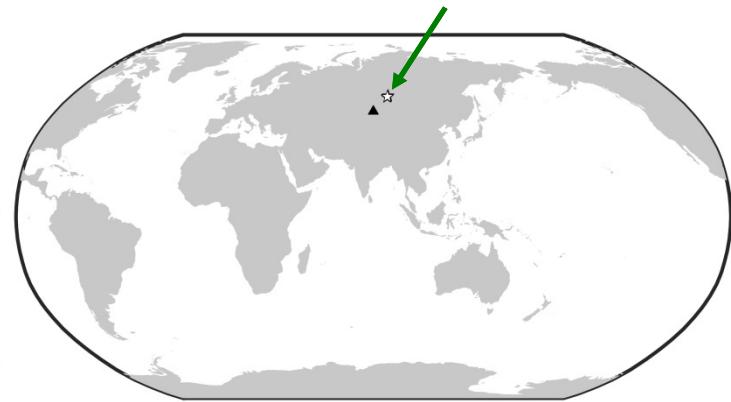
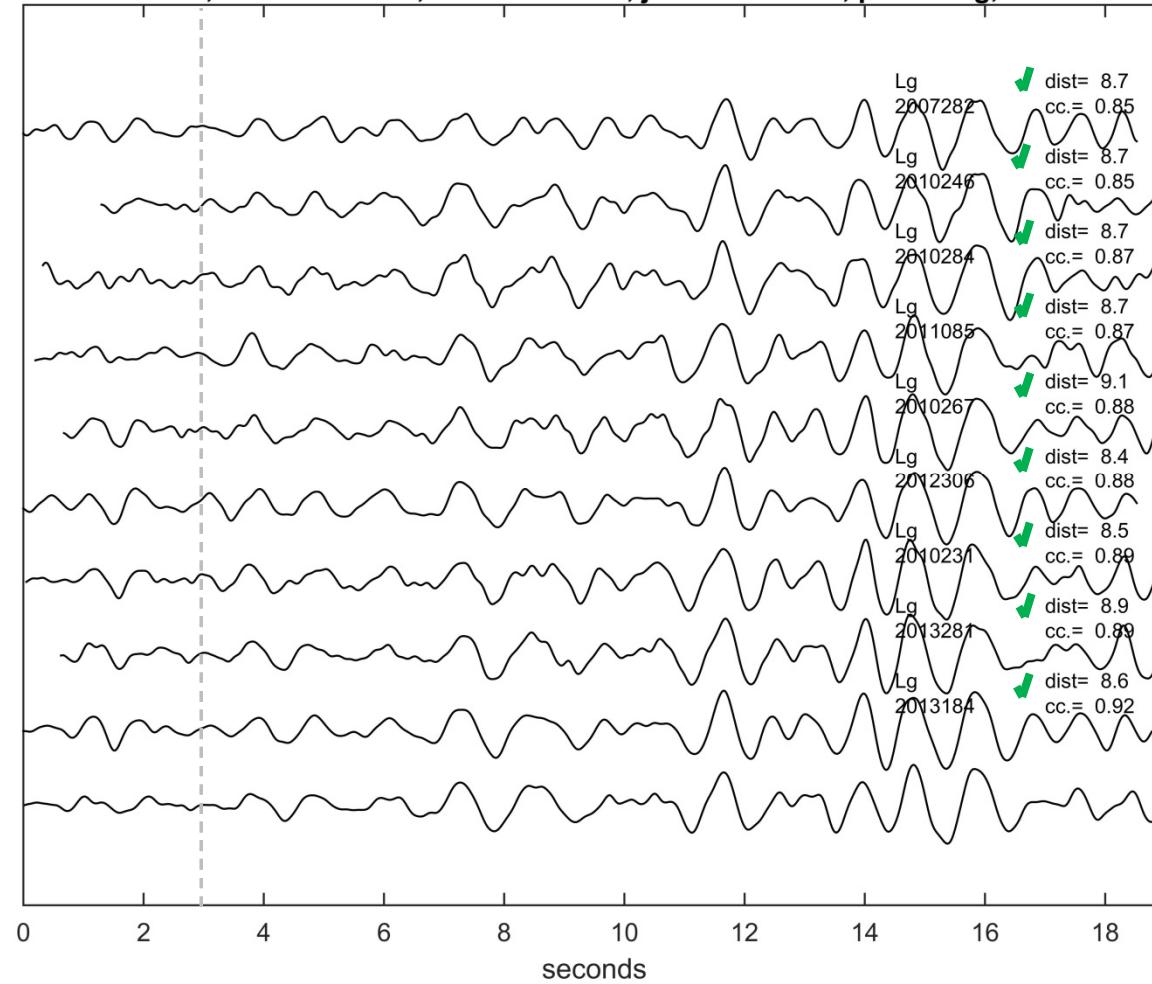


More on Lg Example: Beware of “matches” with low correlation scores



Regional Accepted Match Example #1 (Lg)

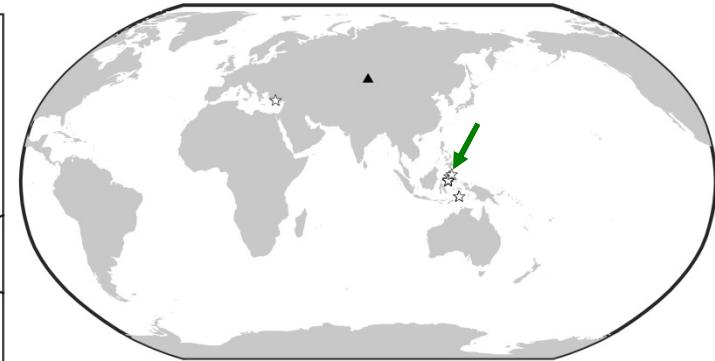
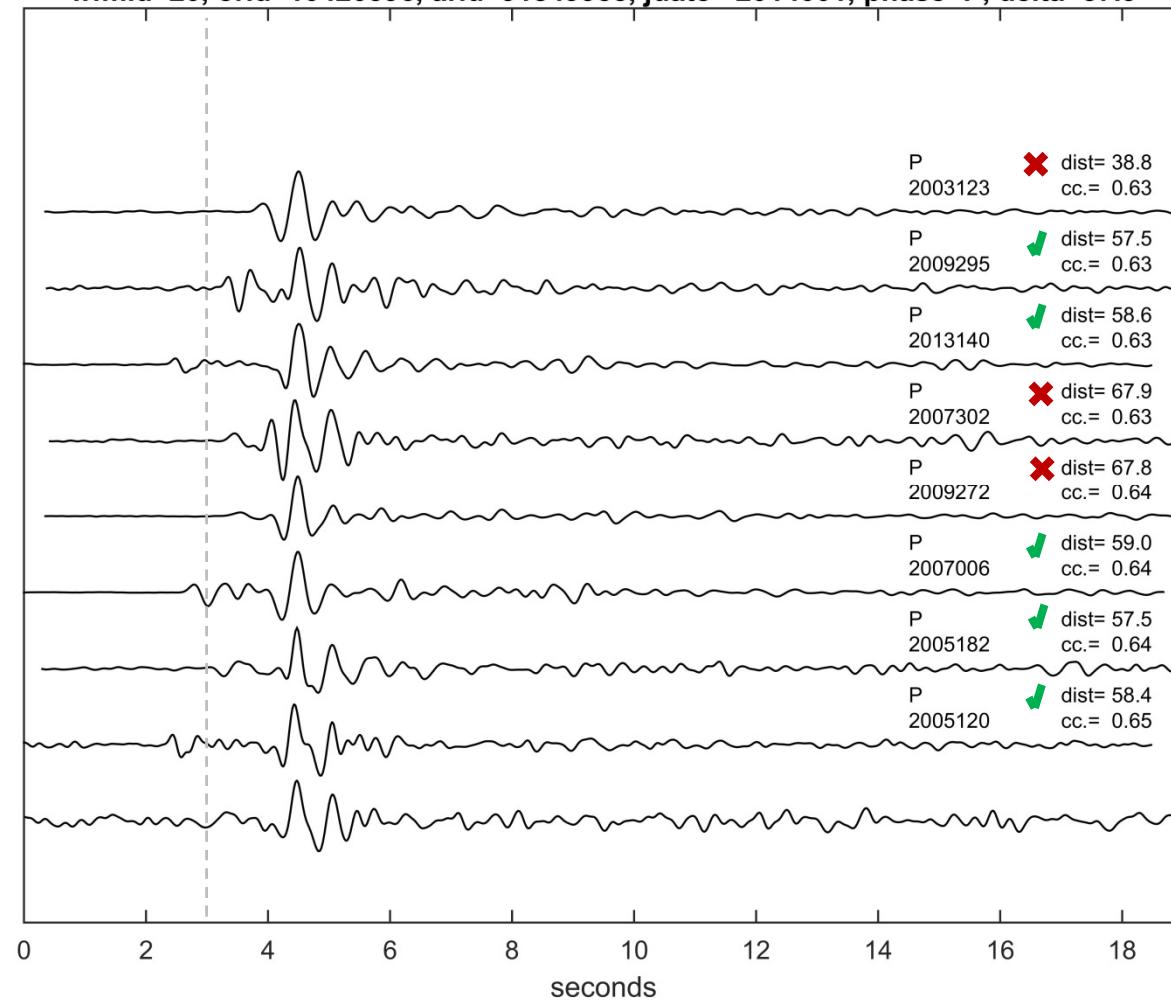
wfmid=248, orid=10426351, arid=91922945, jdate= 2014004, phase=Lg, delta=8.6



- Complex, long-duration signal, hence no disagreement in match pool
- Match pool span 2007-2013
- Mining explosions?

Teleseismic P Accepted Match Example #2

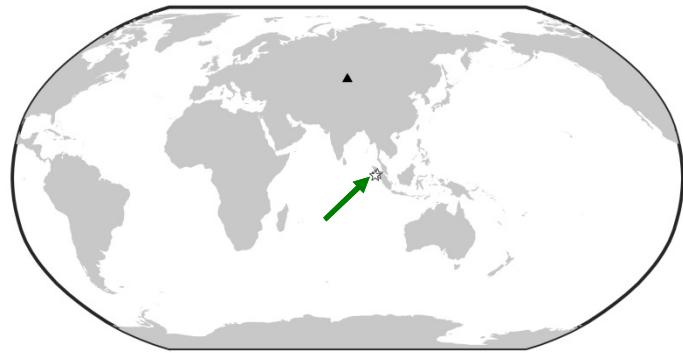
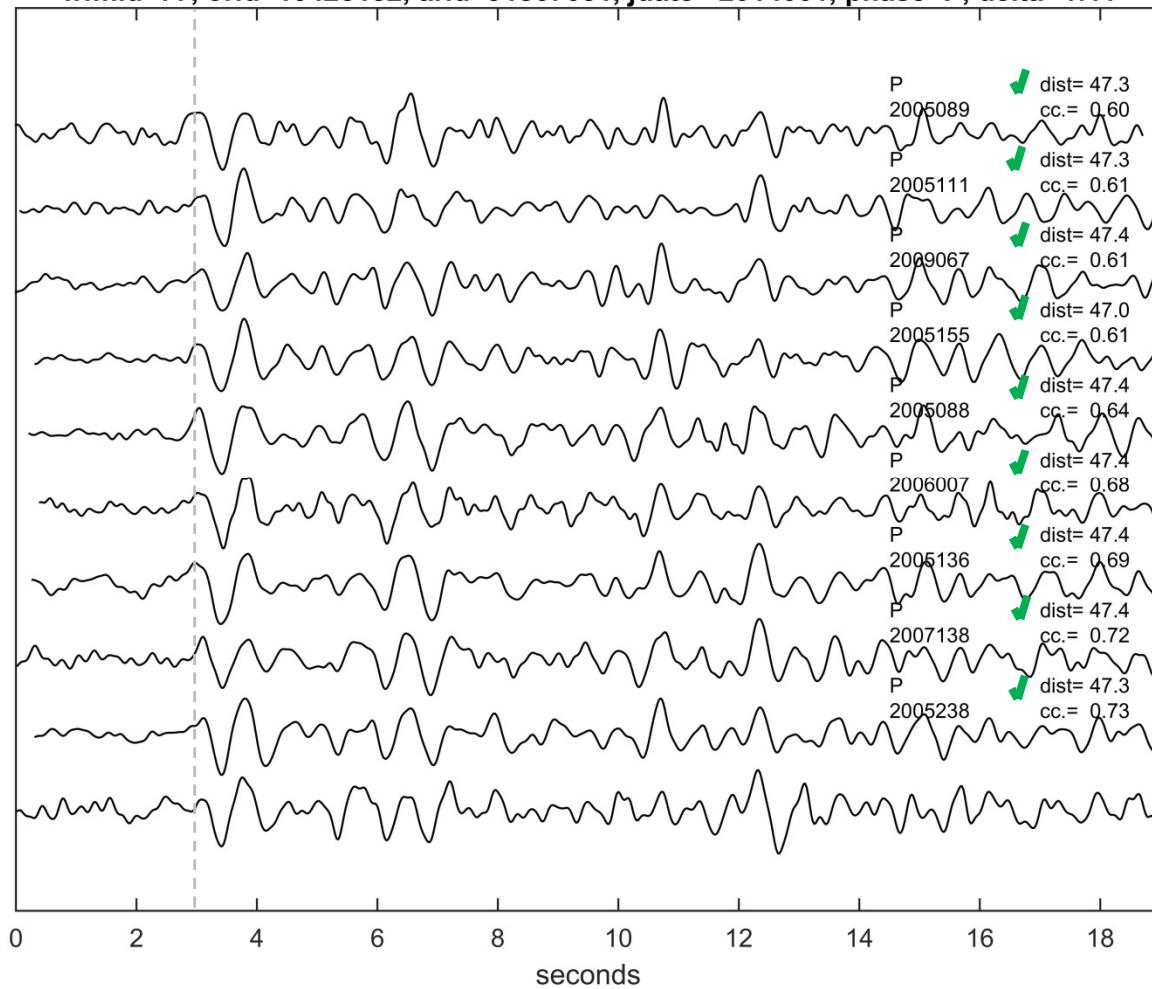
wfmid=20, orid=10420396, arid=91849955, jdate= 2014001, phase=P, delta=57.9



- Pool of matches again small
- 3 of 8 wrong, but still enough consistency in top 4 to accept the match

Teleseismic P Accepted Match Example #3

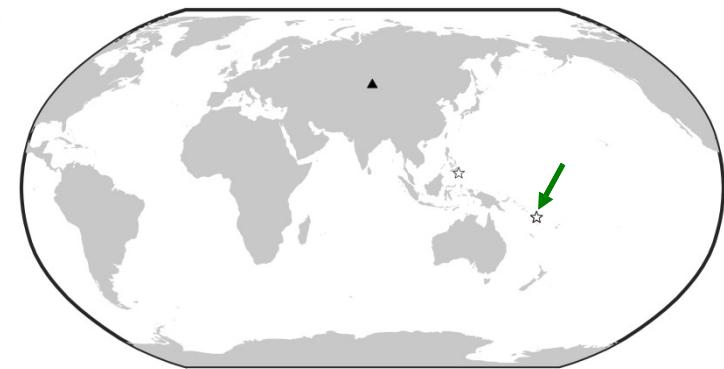
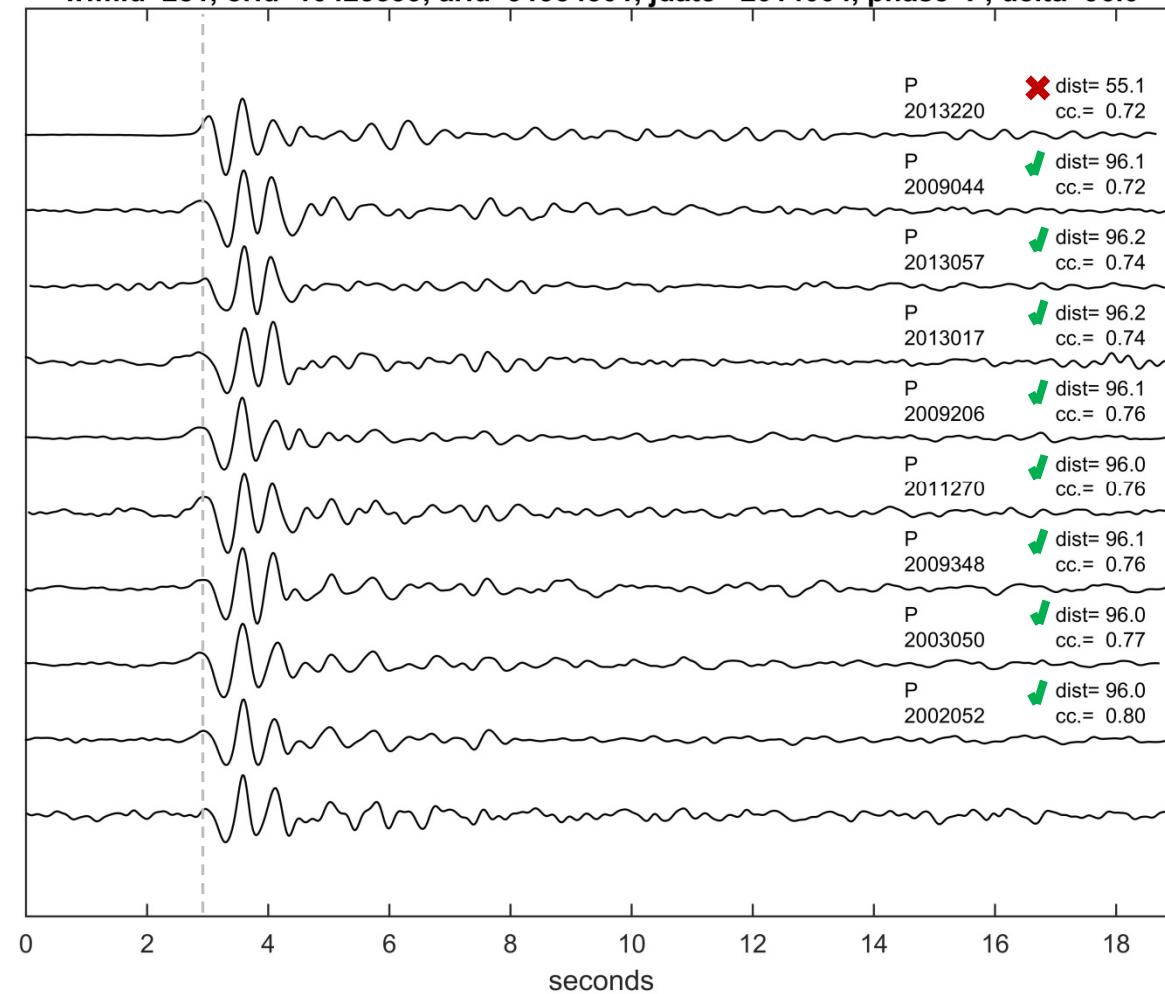
wfmid=77, orid=10425132, arid=91867031, jdate= 2014001, phase=P, delta=47.4



- Long duration, complex signal (path effect?) so perfect agreement in match pool

Teleseismic P Accepted Match Example #4

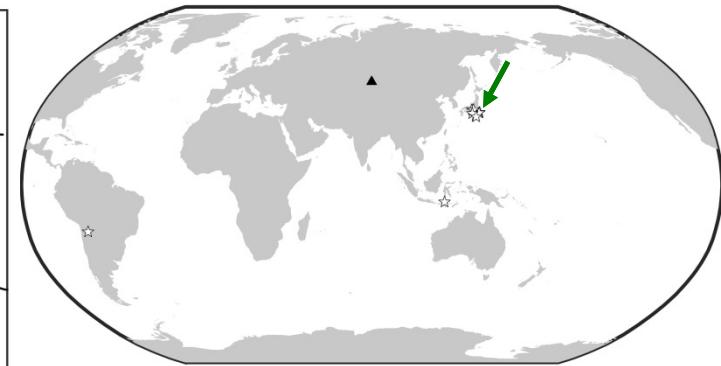
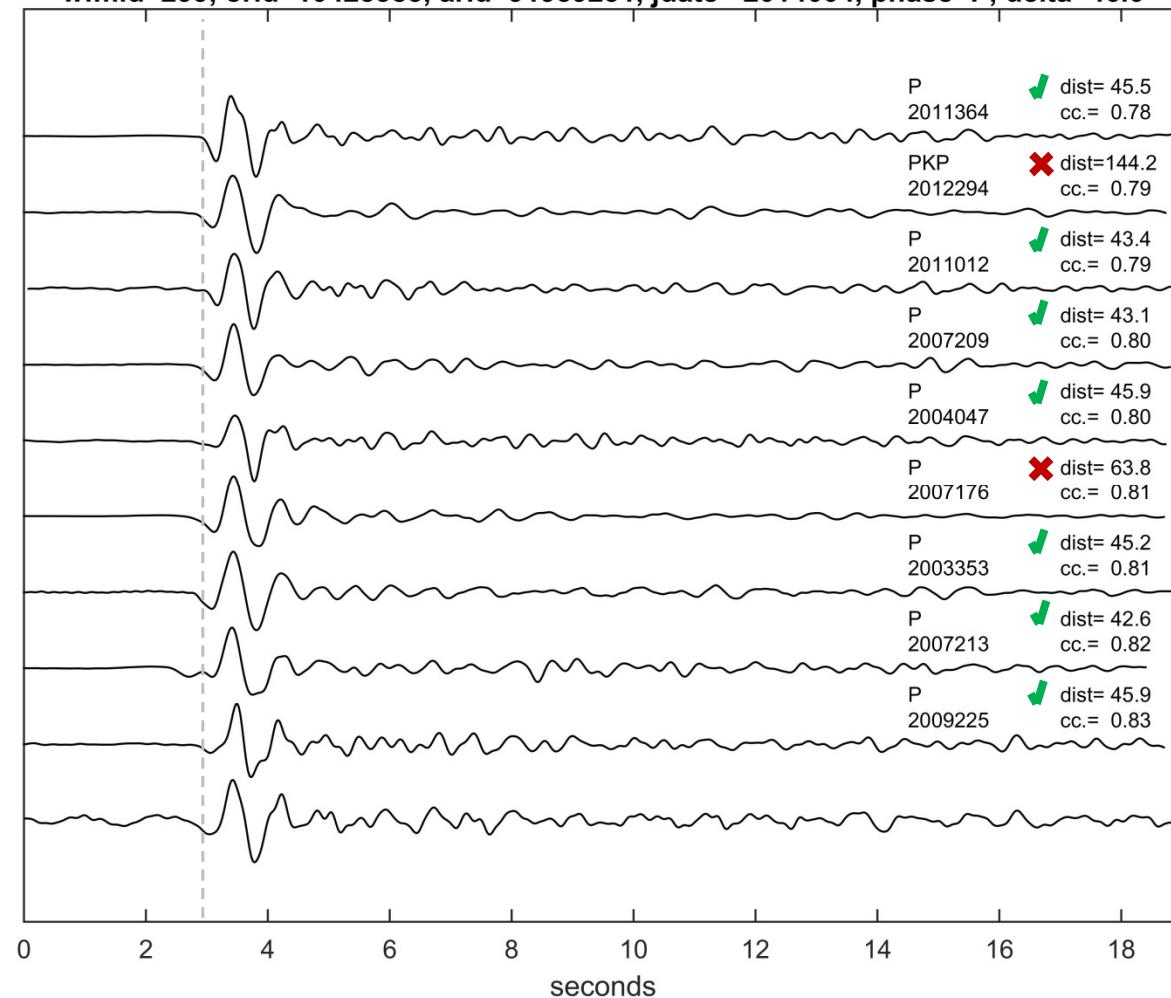
wfmid=281, orid=10423599, arid=91934301, jdate= 2014004, phase=P, delta=96.0



- Moderately complex signal, so good agreement among ANN set

Teleseismic P Accepted Match Example #5

wfmid=299, orid=10425383, arid=91939231, jdate= 2014004, phase=P, delta=46.0



- Signal seems neither long nor very complex, but match is accepted and is correct