Ⓔ

Ocean-Bottom Seismometer Instrument Orientations via Automated Rayleigh-Wave Arrival-Angle Measurements

by Adrian K. Doran and Gabi Laske

Abstract   After more than 10 years of the U.S. ocean-bottom seismometer (OBS) Instrument Pool operations, there is still need for a consistent and accurate procedure to determine the orientation of the horizontal seismometer components of passive-source free-fall broadband OBSs with respect to geographic north. We present a new Python-based, automated, and high-accuracy algorithm to obtain this information during postprocessing of the data. As with some previous methods, our new method Doran–Laske-Orientation-Python (DLOPy) is based on measuring intermediate-period surface-wave arrival angles from teleseismic earthquakes. A crucial new aspect of DLOPy is the consultation of modern global dispersion maps when setting up the analysis window. We repeat measurements at several frequencies to lower biases from wave propagation in laterally heterogeneous structure. We include measurements from the first minor and major great-circle arcs to further lower biases caused by uneven geographical data coverage. We demonstrate the high accuracy of our technique through benchmark tests against a well-established "hands-on" but slow technique using data from instruments of the Global Seismographic Network for which orientations are well documented. We present results for all Cascadia Initiative deployments, along with a number of other OBS experiments. Compared to other widely used automated codes, DLOPy requires fewer events to achieve the same or better accuracy. This advantage may be greatly beneficial for OBS deployments that last as short as a few months. Our computer code is available for download. It requires minimal user input and is optimized to work with data disseminated through the Incorporated Research Institutions for Seismology Data Management Center.

Electronic Supplement: Tables containing station parameters, including location and depth, as well as calculated orientation for the H1 component of the broadband sensors.

Introduction

Many if not most seismic applications that use horizontal seismometer components are crucially dependent on the accurate knowledge of the instrument orientation, that is, the alignment of the nominal 1 and 2 components with geographic north and east. Seismometers on land can be oriented quite accurately, for example, to within a fraction of a degree using a gyro (see, e.g., Ekström and Busby, 2008; Ringler et al., 2013, for reviews). Network operators then either align the components with the geographic coordinates or report a misalignment in the corresponding metadata that accompany a data stream. A data user can take this information into account and rotate the horizontal components accordingly.

The determination of the orientation of free-fall ocean-bottom seismometers (OBSs) is a much greater challenge.

Attempts to use a fluxgate magnetometer to gauge the component orientation upon arrival on the seafloor are subject to large errors (J. Collins, personal comm., 2012). In active-source experiments, the instrument orientation can be determined through airgun shots from several locations (e.g., Anderson et al., 1987; Duennebier et al., 1987; Riedel et al., 2014), but this practice is not common for most of the recent passive broadband OBS deployments because of ship-time costs and permitting challenges. The data user, which includes the principal investigator of a specific experiment, often has to determine the component orientation during postprocessing. Very often, this is done concurrently with the analysis of seismic observables, such as shear-wave splitting or receiver functions (RFs). Accurate determination of

instrument orientations requires global azimuthal coverage to remove biasing effects from 3D structure, a requirement that is often violated (Laske et al., 1994). Instrument orientations obtained during postprocessing may be documented in publications (e.g., Stachnik et al., 2012; Rychert et al., 2013; Zha et al., 2013; Janiszewski and Abers, 2015), but these results are rarely reported to the Incorporated Research Institutions for Seismology Data Management Center (IRIS-DMC) and disseminated as metadata to the seismic datasets. In addition, no method has emerged as a standardized and widely accepted practice, leading to inconsistencies in reported OBS orientations.

A number of numerical approaches and seismic techniques exist to obtain instrument orientations. Some investigators exploit the particle motion of P waves, either by examining the first-arrival particle motion (e.g., Yoshizawa et al., 1999; Schulte-Pelkum et al., 2001) or by minimizing the energy on the transverse component in a selected window (e.g., Niu and Li, 2011). Janiszewski and Abers (2015) describe a method to orient horizontal OBS components by minimizing the RF energy on the transverse component. Zha et al. (2013) use the interstation Rayleigh-wave impulse response recovered from ambient noise correlations. Ekström and Busby (2008) generate synthetic seismograms using previously calculated moment tensor solutions and find the arrival angle that best correlates with the surface-wave time series. Stachnik et al. (2012) developed an automated procedure to measure arrival angles of first-arriving Rayleigh wavetrains (R_1), averaged over a wide frequency band. This algorithm is currently used by the IRIS OBS Management Office (IRIS OMO) to provide initial OBS orientation estimates for OBS community experiments such as the Cascadia Initiative (CI; Toomey et al., 2014). The OBS Instrument Pool (OBSIP) website offers a corresponding MATLAB toolbox (see Data and Resources) under the data tag for use by others (version date as downloaded on 14 November 2014) and is described in Sumy et al. (2015). For comparison purposes with our new technique, we wrote a Python version of this toolbox and verified its consistency with the MATLAB toolbox as well as the strategy of the original Stachnik et al. (2012) publication, except for the way the MATLAB toolbox computes error bars. In the following, we will refer to this Python implementation as the Stachnik et al. (2012) method or STACHPy. One of the first techniques to determine instrument orientation at Global Seismographic Network (GSN) stations during postprocessing was the "hands-on" approach of Laske et al. (1994), who measured surface-wave arrival angles interactively at a range of frequencies, for both Rayleigh and Love waves, for the first-arriving but also later wavetrains. This Fortran and C-based interactive-screen tool is quite accurate but considered too slow and labor intensive for many modern networks with large numbers of stations, on one hand, and the need for rapid, community-service-type determinations of instrument orientations on the other.

Here, we introduce a Python- and Fortran-based hybrid, Doran–Laske-Orientation-Python (DLOPy) automated method

that takes advantage of both approaches, that is, the automated Stachnik et al. (2012) R_1-only method but also some features of the interactive but slow Laske et al. (1994) method. With regard to datasets used, we benchmark results using GSN stations but also several legacy OBS deployments such as the Ocean Seismic Network (OSN) pilot deployment (Collins et al., 2001). The ultimate goal of this article is to report an internally consistent set of instrument orientations for all four 1-year OBS deployments for the CI. The CI was the first amphibious (onshore/offshore) community experiment. Seismic data were not subject to the usual 2-year proprietary hold but were made available to all users as soon as the OBS operators uploaded the data to the IRIS-DMC. New instruments were built specifically for the CI by the three OBSIP Institutional Instrument Contributors (IICs): Lamont–Doherty Earth Observatory (LDEO), Scripps Institution of Oceanography (SIO), and Woods Hole Oceanographic Institute (WHOI). As a first for OBS engineering, two groups (LDEO and SIO) developed shielded, trawl-resistant instruments to allow noise-reduced deployment in shallow water on the continental shelf. We find that the SIO Abalone instruments provide consistently high-quality results at water depths greater than 200 m. The advantage of this design seems to hold in deeper waters off Cascadia beyond water depths of 2500 m, though a few nonshielded instruments also yielded high-quality results.

Method

The basic idea behind DLOPy is to measure broadband surface-wave arrival angles for individual earthquakes. The statistical average over all measurements is then taken as the misalignment of the horizontal seismometer components with respect to the geographic coordinate system. We assume that the horizontal components are orthogonal and that the sensor is level. On a heterogeneous Earth, and with uneven source coverage, the statistical average most likely has an often ignored contribution from lateral refraction. We will cover this aspect in the Discussion section.

With respect to channel names, we use the Standard for Exchange of Earthquake Data (SEED)/GSN naming convention: the H1, H2, and Z components form a left-handed coordinate system, with H1 ideally being horizontal component N, H2 being horizontal component E, and Z being up (e.g., Ahern et al., 2012). The H1 and H2 naming convention is used instead of N and E when the horizontal components are misaligned by more than a few degrees (Fig. 1). In many OBS deployments, the components are named using a right-handed coordinate system, with component H1 being the E component. We will explicitly state when a right-handed convention is used.

Our automated method builds upon the approach of Stachnik et al. (2012), which in turn is based on the principles of previous back-azimuth estimation studies (e.g., Chael, 1997; Selby, 2001; Baker and Stevens, 2004). The basic idea is that an isolated Rayleigh wave in a three-component
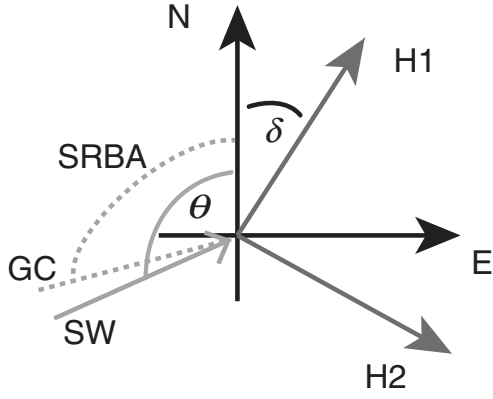
Figure 1. Orientation of the horizontal seismometer components (N and E) using the Global Seismographic Network (GSN)/Standard for Exchange of Earthquake Data (SEED) naming convention for a left-handed coordinate system. Component Z emerges out of the page toward the reader. Misaligned components are named H1 and H2 accordingly. The angle δ is the instrument orientation. Angle $\alpha = \theta + \delta$ is determined by grid search (see the Method section for details). The measured arrival angle as used for further processing is $\alpha -$SRBA (source–receiver back azimuth). In a heterogeneous Earth, the actual approach of a surface wave (SW) may deviate from the source–receiver great-circle (GC) by several degrees.

seismogram, which is rotated into a ray-based coordinate system, appears only on the radial component but not on the transverse. The radial component then correlates with the Hilbert-transformed vertical component. In a grid search, we rotate the horizontal components by an angle α, with

$$\begin{bmatrix} H_1' \\ H_2' \end{bmatrix} = \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix} \begin{bmatrix} H_1 \\ H_2 \end{bmatrix}, \qquad (1)$$

in which α varies between 0° and 360°, at 0.25° intervals. We search for α that maximizes the cross correlation between the Hilbert-transformed vertical component and H_1'. In this case, we find $\alpha = \theta + \delta$, as defined in Figure 1, and H_1' is the radial component. After correction for the source–receiver back azimuth, this angle is the measured arrival angle. As proposed by Stachnik *et al.* (2012), we use the unbounded cross correlation C_{zr}^* to determine the arrival angle because it has clear maxima and minima, but use the normalized correlation C_{zr} to determine the quality of the measurement because the value is bounded on the interval $[-1, 1]$. The relevant functions are defined as

$$C_{zr} = \frac{S_{zr}}{\sqrt{S_{zz}S_{rr}}} \quad C_{zr}^* = \frac{S_{zr}}{S_{zz}} \quad S_{ij} = \sum_{\tau=1}^{N} x_i(\tau)x_j(\tau), \quad (2)$$

in which S_{ij} is the zero-lag cross-correlation coefficient (we also use the term cross-correlation value) between two time series x_i (the radial component H_1') and x_j (the Hilbert-transformed vertical component Z). A user can choose a threshold for the cross correlation below which a measurement is discarded.

Both methods discussed here use a 10% Tukey window as a data taper (Harris, 1978). The most fundamental differ-

ence between DLOPy and STACHPy is the timing of the data taper used for analysis in combination with the protocol of band-pass filtering the wavepackets. Starting with the latter, for a given source–receiver R_1 wavepacket, STACHPy makes a single measurement for a wavepacket filtered between 20 and 40 mHz. The choice of a relatively short window at a fixed length of 10 min 20 s provides uniform spectral smoothing for an entire arrival-angle database, though it does not take into account the different dispersive properties in the wavepackets that change with varying epicentral distances. We choose to make individual measurements on wavepackets that were band-pass filtered around a suite of frequencies between 10 and 40 mHz, at 5 mHz intervals. Associated four-pole, zero-phase Butterworth filters have corner frequencies at ± 5 mHz around the target frequency. The individual measurements enter the final averaging process for the retrieval of the instrument orientation with equal weight, which we will discuss later. The reason why we include lower frequencies than Stachnik *et al.* (2012) is that low-frequency arrival angles are usually less affected by wave propagation in the heterogeneous Earth. In fact, global broadband studies use frequencies as low as 5 mHz (e.g., Laske *et al.*, 1994; Larson and Ekström, 2002). In ocean environments, high noise levels from infragravity waves inhibit useful analysis at frequencies much below 10 mHz (Webb, 1998), and noise levels are particularly high at shallow sites (Webb and Crawford, 2010). On the other hand, our experience has been that many deep-ocean OBSs provide high-quality and consistent dispersion data to frequencies as low as 10 mHz. The background noise on OBSs is depth dependent (e.g., Webb and Crawford, 1999; Yang *et al.*, 2012), and we will investigate the relationship between orientation uncertainty and ocean depth later in this article.

The most crucial difference between our methods is the timing of the data taper. STACHPy uses a fixed R_1 time window that always starts 20 s before a predicted 4.0-km/s phase arrival that varies with epicentral distance. The taper extends to 600 s after this arrival (Fig. 2), regardless of epicentral distance. In contrast, the timing of our tapers depends on frequency and is determined using modern global dispersion maps (Ma and Masters, 2014; Ma *et al.*, 2014) that are sampled in equal-area 1° cells. For each source–receiver pair, we integrate the group travel time to determine the center time of our taper. The determination of this time is currently the only piece of code written in Fortran (courtesy of Zhitu Ma), whereas the rest is in Python. The taper length varies with frequency and was determined somewhat empirically. We started out using a length of 600 s, similar to what Stachnik *et al.* (2012) did, but we subsequently varied the length to optimize the cross correlation between the radial and Hilbert-transformed vertical components. For each frequency, we then picked the length that gave the best correlation for most epicentral-distance ranges. Taper lengths range from 500 s at 40 mHz to 700 s at 10 mHz, which means that the taper for 40 mHz is 20 periods long but at 10 mHz is only seven.
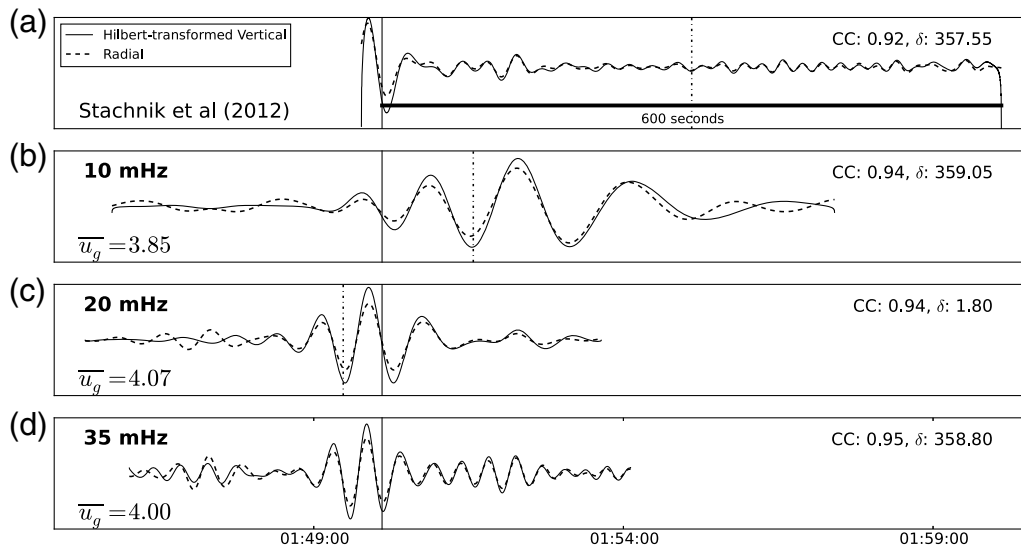
Figure 2. R_1 wavepackets for the 6 February 2013 M_w 8.0 Santa Cruz Islands earthquake (01:12:27 UTC, source depth 24 km) recorded at Cascadia Initiative (CI) station G20B at epicentral distance $\Delta = 81°$. The wavepackets are shown as filtered and windowed but not yet tapered for analysis by (a) Stachnik *et al.* (2012) method (STACHPy), and by (b–d) Doran–Laske-Orientation-Python (DLOPy) at three frequencies. Each panel shows the match between the Hilbert-transformed vertical trace (solid) and the optimally rotated BH1 trace (radial; dashed). Corresponding correlation coefficients (CC) and the inferred instrument orientation for this particular event (δ) are shown in the upper right corner. The traces are amplitude-normalized in each panel for optimal display. The solid vertical bar marks the arrival time for a 4 km/s default group velocity u_g chosen by Stachnik *et al.* (2012), whereas the dotted lines in (b–d) mark the center of our taper. Associated path-averaged source–receiver group velocities are given in the left corner of each panel. The dotted line in (a) marks the center of the STACHPy window and corresponds to a path-averaged group velocity of 3.53 km/s. This window nearly completely misses R_1, which decreases the likelihood for a high-quality arrival-angle estimate.

These values are currently embedded in our code but could be changed by an experienced user.

An example of the dramatic difference in the application of tapers between DLOPy and STACHPy is shown in Figure 2. Although Stachnik *et al.* (2012) do not analyze data at 10 mHz, the associated wavepacket would be the only one considered sufficiently well windowed for analysis, whereas the waveforms for 20 and 35 mHz that are analyzed using STACHPy are nearly missed completely because the window begins too late. We can only speculate that the default velocity of 4 km/s was chosen upon inspecting group velocity curves for reference Earth models, such as preliminary reference Earth model, (PREM; Dziewonski and Anderson, 1981) to determine the beginning of the time window (-20 s). For such a model, the group velocity never reaches 4 km/s for frequencies between 20 and 40 mHz. For PREM, the velocities range between 3.60 km/s at 40 mHz and 3.89 km/s at 20 mHz (frequencies associated with a group velocity of 4 km/s are around 2.9 mHz). The group velocity associated with the center of a Stachnik window (as opposed to the beginning) depends on epicentral distance. In the case shown in Figure 2, it is 3.53 km/s and still seems reasonable for a PREM velocity. However, oceanic paths have particularly high path-averaged group velocities as is documented by the timing of our new windows, and PREM values are poor representers. We should note that the example shown may be an extreme case, but most paths associated with OBS deployments most likely are significantly oceanic. We

should also note that an experienced user may code a more fitting group velocity. However, this would presuppose that the user inspects the data, which is not typically done using automated tools. Surprisingly enough, we obtain reasonable agreement between individual arrival angles measured with both techniques, as discussed below. It is interesting to note that the maximum in the envelope of our wavepackets lag our predicted arrival times for 10 mHz but leads for 20 mHz. This indicates some small inconsistencies between our data and the modern group velocity maps of Ma *et al.* (2014). However, the study of this is not the topic of this article.

Rounding up the basic differences between STACHPy and DLOPy, the former analyzes only R_1. We include the analysis of the major arc R_2, as was done in our previous global studies (e.g., Laske and Masters, 1996). Surface waves experience different lateral refraction on the minor-arc and major-arc great-circle paths and therefore provide independent arrival-angle estimates. Quite often, R_2 has a superb signal-to-noise ratio (SNR), particularly for the larger earthquakes. R_2 is also often cleaner than R_1, in the sense that we get higher correlation between the radial and the Hilbert-transformed vertical components. We also observe higher singular values with the older Laske *et al.* (1994) method, which indicates that single, isolated signals approach the station (as opposed to multiple, interfering signals). R_1 can be contaminated by coda signals from the Love wave G_1 as well as from Rayleigh-wave overtones. Because of different dispersion, their effects on R_2 are different, so that R_2 may be more isolated.

Implementation and Application

For our analyses, we use all global shallow earthquakes with M_w or $M_s \geq 5.5$. We exclude events with source depths greater than 150 km to lower the risk of overtone contamination. We also exclude events with epicentral distances $\Delta < 5°$ and $\Delta > 175°$ to avoid bias from near-source and antipodal multipathing effects. This choice also ensures that all sources are at least one wavelength away from the station for our longest-period measurements. Our computer code utilizes the recently developed seismic package ObsPy (Beyreuther *et al.,* 2010; Megies *et al.,* 2011; Krischer *et al.,* 2015) to download data from the IRIS-DMC. The code is written primarily using Python 2.7 (as opposed to Python 3.0 or later releases) to simplify user interaction as much as possible. For each station separately, we determine the final station orientation and uncertainty using the following protocol:

1. retain all measurements with $C_{zr} \geq 0.80$ (adjusted as needed depending on noise conditions at individual instrument locations);
2. remove outliers by computing the median of data and retaining all values within five times the median absolute deviation (MAD);
3. use the bootstrap method of random sampling with replacement (Efron, 1979) to compute the arithmetic mean after 5000 runs on the remaining values and the 95% confidence interval of the mean. We report the uncertainty as twice this confidence interval.

In step 2, using the MAD to identify outliers is preferable to methods involving the standard deviation or variance of the data, because these statistics yield biased results for datasets with severe outliers. As datasets are resampled during bootstrapping, this strategy may allow us to remove some of the biasing effects caused by the uneven azimuthal distribution of earthquakes. A MAD cutoff of 5 was determined empirically. It may seem somewhat conservative but produces robust results.

Step 3 is a new way to determine error bars that was not implemented by Stachnik *et al.* (2012). Because the error of the mean is inversely proportional to the square root of the number of calculations (standard error $\approx \sigma / \sqrt{n}$), more calculations generally lead to a more precise estimate, if the underlying dataset has no systematic biases. Although many investigators present orientation uncertainties using the 95% confidence interval of the mean, we choose twice this value, effectively reporting the 4-σ uncertainty. As will be discussed below, this more conservative choice accounts for some systematic biases in the datasets. It also ensures internal consistency and repeatability within the dataset as well as more consistency when using different measurement strategies. We should add that because we use the bootstrap method to determine error bars, repeat runs yield different error bars. Our experience is that errors may change by 0.05° for stations having high SNRs, which is actually a very small number compared to the actual error bars.

As an example, Figure 3 displays an entire arrival-angle dataset for CI station G20B. Obviously, DLOPy gives more measurements than STACHPy because the latter gives only frequency averages. In both cases, the vast majority of data have low cross-correlation values, which implies low quality. However, high-quality data associated with cross-correlation values higher than a certain threshold tend to cluster around an average, which we take as the station orientation. For a given threshold, our data tend to scatter more than those using STACHPy, most likely because of the implicit spectral averaging in the latter. Hence, the scatter alone does not rank the quality of one method over the other. Quite often, only one or two frequencies produce high-quality results for R_1, so a measurement using STACHPy may produce a low-quality result that is ultimately discarded, whereas they are retained as data using our code. Also recall that the dataset using our code includes many high-quality R_2 measurements.

For CI station G20B station, we obtain an orientation of $2.39 \pm 2.82°$ from 209 high-quality measurements for 76 events. This includes 180 R_1 data and 29 R_2 data. Application of STACHPy yields an orientation of $0.26 \pm 6.08°$ from only 10 events. We should expect that the vastly lower number of events used in the latter bears an increased risk of biasing effects by dominant ray corridors on a laterally heterogeneous Earth. Nevertheless, both our methods yield consistent station orientations, with statistically insignificant discrepancies. However, our error bar is smaller than that obtained using STACHPy for which we did not multiply the error by a factor 2. The IRIS OMO has recently started implementing the Stachnik *et al.* (2012) method to determine the OBS orientations for community experiments such as the CI. OMO used 28 events to obtain an orientation angle of $3.00 \pm 15°$. The IRIS OMO MATLAB implementation computes errors in a different way, namely choosing the standard deviation in the dataset. In the strict statistical sense, this is too conservative on one hand but gives unrealistically small errors for small datasets.

Benchmark Methods and Datasets

In this section, we attempt to validate several aspects of the measurement process. A crucial question is whether relatively short deployments (or a collection of only a few earthquakes) can provide reasonably accurate estimates of instrument orientations. We address this using data from GSN stations ESK (Eskdalemuir, Scotland) and HRV (Harvard, Massachusetts).

We also assess the validity and repeatability of our results, and the automated methods in general. For this, we compare our results for GSN station SACV (Santiago Island, Cape Verde) against those obtained using the interactive Laske *et al.* (1994) benchmark method. In this method, the frequency-dependent surface-wave arrival angles are measured by hand via an interactive screen tool. The data analyst chooses the window that best visually isolates the desired waveforms in the three-component seismogram. A singular value decomposition of the complex three-component spec-
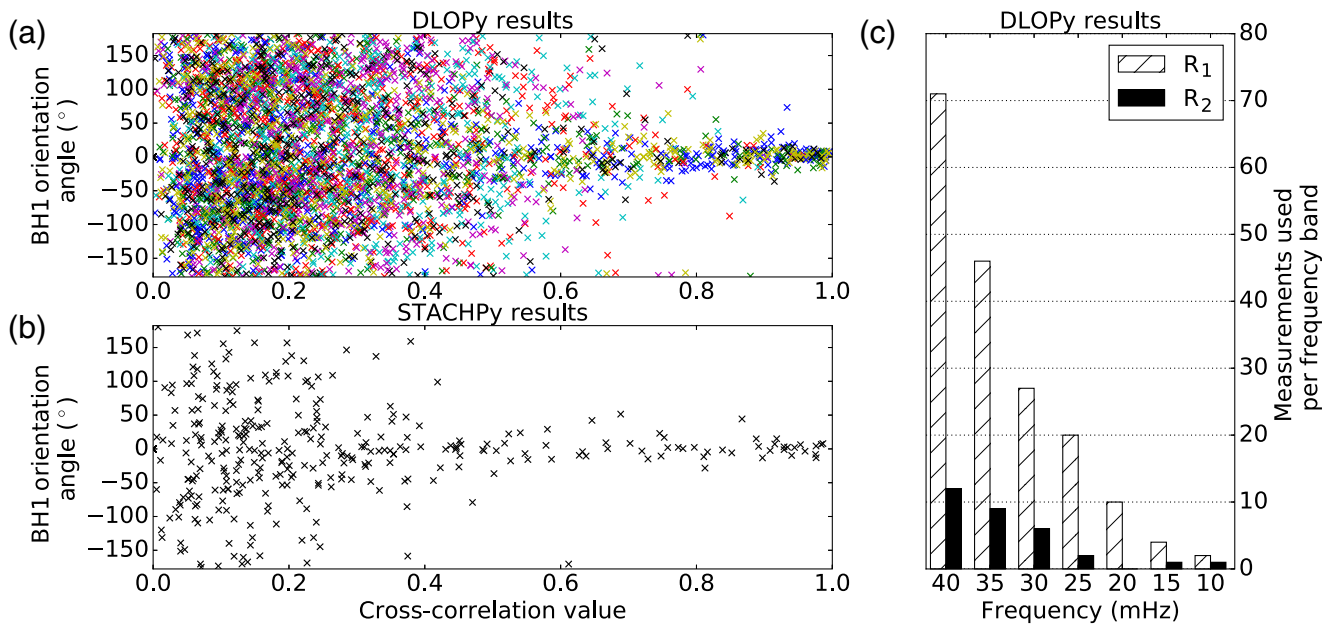
Figure 3. A complete set of arrival-angle measurements as a function of cross-correlation value for CI station G20B. (a) Doran and Laske method DLOPy; (b) implementation of Stachnik *et al.* (2012) method STACHPy. The cross-correlation value is a measure for data quality. "BH1 orientation angle" is the optimum rotation angle α that maximizes cross correlation between component H_1' and the Hilbert-transformed vertical component. Different colors (or shades of gray in the print edition) in (a) mark measurements at different frequencies and wave orbit number (1 or 2). High-quality data tend to cluster around a common value, the inferred instrument orientation δ. Coincidentally, this instrument was almost naturally oriented. (c) A count of the number of high-quality measurements used from each great-circle arc and frequency. At G20B, the greatest number of usable measurements are made for R_1 at 40 mHz. Final instrument orientations: STACHPy, $0.26° \pm 6.08°$; DLOPy, $2.39° \pm 2.82°$; and Incorporated Research Institutions for Seismology Ocean-Bottom Seismometer (OBS) Management Office (IRIS OMO) report, $3.00 \pm 15°$. The color version of this figure is available only in the electronic edition.

trum then yields frequency-dependent eigenvectors that govern the principal particle motion. The associated singular values give information on the quality of the motion. In the ideal case of an isolated signal with particle motion in one plane (including linear motion), one singular value is one and the other two are zero. If needed, the data analyst finds the best window iteratively that yields the highest singular value. A multitaper approach optimizes between spectral leakage effects and bias from noise in the seismic records. The orthogonal multitapers also provide statistically independent estimates, allowing the assignment of formal error bars for each individual measurement. The interactive choice of the data window is particularly effective and important in the case of Love waves to make sure that no signal of the earlier or concurrent Rayleigh-wave overtones nor the later Rayleigh-wave fundamental mode is in the window.

We then proceed with some past temporary and permanent OBS deployments. We compare surface and buried OBSs during the short-term OSN pilot experiment (OSNPE) near Hawaii (Collins *et al.*, 2001). We also investigate the instrument orientation of the buried OBS at the high-quality, long-term MOnterey Bay Broadband Observatory (MOBB; Romanowicz *et al.*, 2006) offshore California. Finally, we benchmark our results for the phase 2 deployment for Plume-Lithosphere Undersea Mantle Experiment (PLUME2) near Hawaii (Laske *et al.*, 2009) against published values obtained with the Laske *et al.* (1994) method.

Internal Consistency and Deployment Length on the Examples of ESK and HRV

We begin our benchmark testing by applying DLOPy at a particular station for different time periods. For this test, we need a station that has had a consistent seismmometer deployment without interruption for a long time, and we choose station ESK, operated by the International Deployment of Accelerometers group with IRIS network code II. According to IRIS-DMC metadata information, this station operated a Wielandt–Streckeisen very broadband (VBB) STS-1 seismmometer without interruption for six years between 1 January 2004 and 1 January 2010. This includes no downtime, no changes to the instrument response, no recovery and redeployment of the sensor, and no other adjustments. The STS-1 is the primary sensor at ESK and carries the location code 00.

For our test, we determine instrument orientations using data for the entire six-year period as well as several subsets of the data spanning one full year, six months, and three months. The results are summarized in Figure 4. As expected, the dataset of the full six years yields the final instrument orientation with the smallest error bar: $3.34° \pm 0.48°$. Results obtained for shorter time spans, for example, one year or several months, yield orientations that slightly vary, but the six-year benchmark result is within the respective error bars. Some of this variation may be due to differing azimuthal distributions of earthquakes during the time periods considered. Larger variations have to be expected for shorter time spans

that implicitly incorporate a smaller number of earthquakes, but the larger error bars make these variations statistically insignificant. This indicates that it is possible to obtain reasonably well the instrument orientation for deployments as short as a few months. However, orientations of short deployments may have error bars of 2.5° or more.

For the epoch spanning the six years, the IRIS-DMC metadata database currently shows an orientation of the N component as 0.1° but the two horizontal components are not orthogonal. The STS-1 components are physically distinct packages, so unintentional nonorthogonal setups are possible though unlikely. The azimuth of the E component is reported as 89.5°, so the angle between the two components is 89.4°. Our code does not account for this. The orientation at ESK is reported at the DMC as being the same since operations started in 1987, so we also compare our results with previously published results. Laske (1995) reported 0.07° ± 0.37°, and her subsequent but unpublished updates in the mid-2000s were as high as 0.28°. These angles were obtained during joint inversions for station orientation and laterally heterogeneous structure. Larson and Ekström (2002) obtained 1.1° ± 0.72° after inversion and 0.9° through simple averaging. We currently cannot explain the relatively large discrepancy with our results of more than 2°. The previous studies included Love-wave angles in their analyses. Our experience is that in some cases this does make a significant difference. We confirmed with the station operator that no realignment of the sensor was performed at ESK (P. Davis, personal comm., 2015), thereby ruling out the possibility that the sensor orientation was changed between studies. Here, we take the 2° difference as a threshold for how accurately and consistently the station orientation can be determined at low-noise stations.

We repeat this procedure at IRIS-U.S. Geological Survey (IU) station HRV to further test the consistency and accuracy of DLOPy. This station also operates a VBB STS-1 seismometer (location code 00), and the DMC reports perfect orthogonality between the horizontal components. The HRV BH channels ran with very few interruptions between 1 January 1991 and 1 January 1997. The IRIS-DMC reports intermittent data gaps between 3–8 December 1996, while construction was occurring in the vault. An instrument response correction affecting December 1996 data was performed retroactively in March 2006. As shown in Figure 4, we see the same consistency in these data as achieved at ESK. The final orientation we obtain is −0.2° ± 0.52°. Our results agree closely with those of Laske (1995) and Larson and Ekström (2002), who report orientations of 0.00° ± 0.41° and −0.2° ± 0.64°, respectively.

Repeatability of Final Instrument Orientation on the Example of SACV

A possible reason for the discrepancy in the instrument orientation for station ESK could be because different workers used earthquakes from different time periods. We therefore repeat a test for the time period between 2011.169 and 2012.202 (13 months) using data collected at station SACV.

This is a time period when no changes to the instrument responses were made and no data gaps occurred. The station has two seismometers for which the primary sensor is a Geotech-Teledyne KS54000 VBB sensor (SEED location code 00), whereas a broadband Güralp CMG-3T seismometer serves as secondary sensor (location code 10). KS54000 is a borehole sensor, and IRIS-DMC reports an emplacement depth of 97 m. At the time of this study, there is some debate on how accurately the sensor orientations are known for this station (P. Davis, personal comm., 2015). This is one reason why we choose this station for our tests.

Applying the automated methods as described above, we determine the instrument orientation for both sensors (Table 1). Using our new approach, the results are based on high-quality measurements from 454 events for sensor 00 and 343 events for the noisier sensor 10. STACHPy yields error bars that are about twice as large as the ones we obtain with DLOPy, but our orientations of 28.75° and 4.04° for both sensors lie within the error bars of the STACHPy results. This indicates once again that at low-noise GSN stations, both methods yield mutually consistent results for a 1-year long deployment.

Using the Laske et al. (1994) interactive-screen method, the data analyst usually discards measurements for a wavetrain if the largest singular value remains below 0.6 or so. Unlike with our new method, those low-quality results never enter the database. For this study, only results with singular values greater than 0.7 are included in further processing. We determine the median of all minor and major arc results at six frequencies: 5, 8, 11, 14, 17, and 20 mHz, for each Love and Rayleigh wave. The average of the measurements gives the final instrument orientation, and the standard deviation divided by 12 gives the uncertainty. We choose to work with the median rather than the weighted averages in this particular test because we do no further vetting against outlier data that we otherwise discard during secondary screening. For the location code 00 sensor, we obtain about 180 Rayleigh-wave high-quality measurements and 160 Love-wave measurements, at each of the six frequencies. The location code 10 sensor is less well constrained, with about 80 Rayleigh-wave and 100 Love-wave measurements. For the two sensors, we obtain instrument orientations of 28.59° and 4.02° (Table 1). Both values are in excellent agreement with those obtained with DLOPy. The Laske et al. (1994) values are now reported as instrument orientations in the metadata at the IRIS-DMC. We note that the DMC reports that the horizontal components are not quite orthogonal, with a 90.4° angle between them.

Because the Laske et al. (1994) method includes Love-wave angles but our new method does not, Table 1 also lists final instrument orientations when excluding Love waves. Although results for the quieter 00 sensor agree with those using DLOPy to within less than a degree, the noisier 10 sensor for which less data are available now has a larger discrepancy of 1.4°. This is still within our error bars though.

To test the repeatability of the obtained instrument orientations, a novice data analyst repeated the Laske et al. (1994) method on the dataset for sensor 00. The analyst required
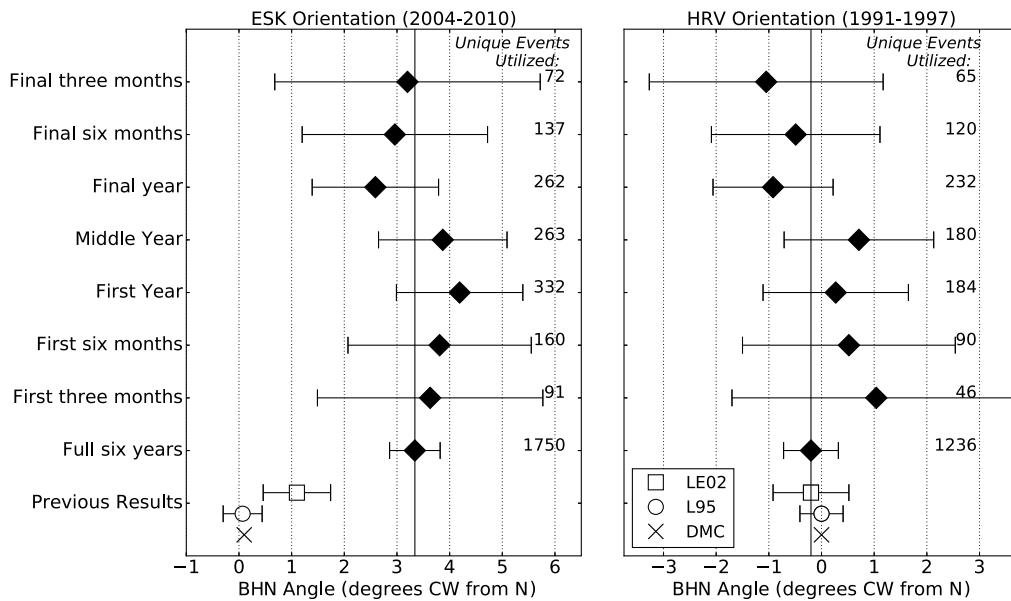
Figure 4. Instrument orientation consistency check at GSN stations ESK (Eskdalemuir, Scotland) and HRV (Harvard, Massachusetts). Results are shown when obtained using the full dataset for six years and several subsets thereof. The benchmark results for the full six-year dataset are $3.34° \pm 0.48°$ for ESK and $-0.2° \pm 0.52°$ for HRV. The number of unique earthquakes contributing to each data subset and the results of previous misorientation analyses are shown: LE02 refers to Larson and Ekström (2002), L95 refers to Laske (1995), and DMC refers to the BHN orientation reported on the IRIS Data Management Center (IRIS-DMC). The orientation of BHN component is shown in degrees clockwise (CW) from geographic north (N).

much more time than Laske to make the measurements, but his efforts yielded an internally more consistent dataset, particularly at low frequencies. The final instrument orientation he obtained is $28.69° \pm 1.66°$ and is in excellent agreement with Laske's result. Addressing the outlier issue, we repeat our procedure after we discarded outliers in Laske's dataset. Histograms reveal that the data distribution is somewhat one-sided, probably as a result of uneven earthquake distribution and wave propagation on a heterogeneous Earth. We therefore choose cutoff limits by hand instead of applying a standard n-σ cutoff but leave the bounds constant across all frequencies. Overall, we discard data that are about 2.5σ or more away from the median at high frequencies but only 1.5σ at low frequencies. The final instrument orientations after this data vetting are $29.43° \pm 0.87°$ and $4.41° \pm 1.01°$. The data-cleaning process clearly yields smaller error bars, by a factor of almost 2, but does not significantly change the final instrument orientations. Nevertheless, these changes are larger than the original 2-σ error bars obtained with our new automated method DLOPy. To accommodate all these discrepancies we choose to work with 4-σ errors instead, as noted above.

Repeatability of Individual Measurements on the Example of SACV

After exploring consistencies and discrepancies in the final estimates of instrument orientation, we now investigate the consistency of each individual measurement. This check is particularly important for the use of surface-wave arrival angles in interpretations of lateral heterogeneity (e.g., Laske

and Masters, 1996, 1998; Larson *et al.*, 1998; Yoshizawa *et al.*, 1999; Larson and Ekström, 2002). A quick comparison between the individual measurements taken by Laske using the Laske *et al.* (1994) method agree reasonably well with those taken by the novice data analyst, in that most measurements lie within their respective error bars. Some outliers include the assignment of the wrong wave orbit or wave type and likely would not pass secondary screening. The interesting and important question is how the Laske *et al.* (1994) dataset compares with the one obtained with our new method. This is demonstrated in Figure 5.

Here, we use different selection thresholds to consider only high-quality data. For the Laske *et al.* (1994) method, we increased the singular-value threshold to 0.85 for R_2, whereas we increase the C_{rz} value for DLOPy to 0.85 for R_1 to include only high-quality data but lower it to 0.7 for R_2 angles to allow for more data to be considered. Uncertainties are only shown for Laske *et al.* (1994) because only this technique provides formal error bars for individual measurements. Overall, our measurements are in close agreement, with most of our data placing within the error bars obtained with the Laske *et al.* (1994) method. As discussed above, R_2 wavepackets sometimes appear to be cleaner than R_1, resulting in measurements with smaller errors, and the agreement between our methods supports this impression (by the higher correlation coefficients between the two datasets).

However, we also note some discrepancies. The different windowing of the data may explain some of this. A screen-interactive method allows the analyst to select slightly

Table 1
Comparison of Orientation Results at Global
Seismographic Network station SACV

Method	SACV.00 (°)*	SACV.10 (°)*
DLOPy	28.75 ± 0.96	4.04 ± 1.28
STACHPy	30.05 ± 2.10	4.48 ± 2.04
Laske *et al.* (1994)	28.59 ± 1.67	4.02 ± 1.99
Laske *et al.* (1994) $R_1 + R_2$	29.38 ± 3.34	5.42 ± 3.94

*Results for time period 2011.169 through 2012.202 as obtained with the three methods. We consider both the primary sensor (location code 00) and secondary sensor (location code 10). DLOPy, Doran–Laske-Orientation-Python; STACHPy, the Stachnik *et al.* (2012) method.

different window positions and lengths to emphasize the wavepacket of interest. Although our new automated approach centers the window around a group arrival time using modern dispersion maps, that window may still include some other signals. Unwanted signals, such as the earlier-arriving overtones may then still be included in our new analysis, whereas the Laske *et al.* (1994) windowing explicitly allows an analyst to exclude them. Similarly, multipathing can affect late-arriving energy and results obtained using our new approach. On the other hand, a data analyst using the Laske *et al.* (1994) approach may cut too much into the wavepacket, thereby inadvertently eliminating low-amplitude signals. Our experience is that measurements at low frequencies that are not considered here (e.g., 5 mHz) are particularly vulnerable. On the other hand, because the Laske *et al.* (1994) method determines arrival angles at all frequencies at once, a long window is often required (up to 30 min) to ensure that all of the dispersed energy is captured. This may be too long for the highest frequencies considered (e.g., 20 mHz), and so the Laske data at 20 mHz may be less accurate. Another difference between the methods is the set of tapers used. At the least, the multitapers treat noise in a different way than the single taper we use here. By implementing the same set of multitapers as used in Laske *et al.* (1994), we may achieve better agreement. Finally, the underlying method to determine the optimal angle may also play a significant role. Although Laske *et al.* (1994) perform a formal inversion of the complex spectra of the three-component seismogram, we perform a grid search in the time domain using the cross correlation of two components.

Short- and Long-Term OBS Observatory Sites

We now move from working with land seismometers to those on the ocean floor. We start our investigations at two observatory-quality installations, first at the short-term OSN pilot deployment near Hawaii (Collins *et al.*, 2001; Stephen *et al.*, 2003) and then at the permanent cabled observatory MOBB offshore California (Romanowicz *et al.*, 2006).

OSNPE was operated for four months between February and June 1998 and provided a unique dataset for comparison. For this experiment, three different OBSs were deployed: site OSN1 had a Geotech-Teledyne KS54000 VBB borehole seis-
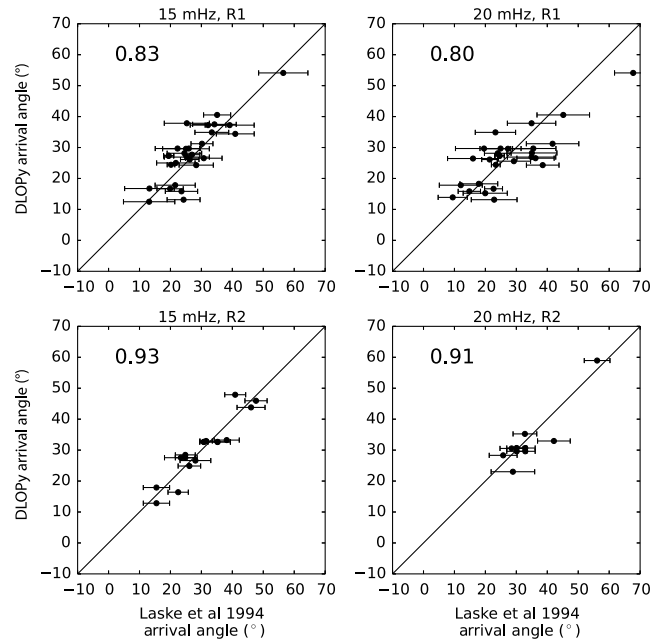

Figure 5. A comparison of individual measurements for GSN station SACV location 00 using the Laske *et al.* (1994) method and DLOPy. Results are shown for two frequencies, 15 and 20 mHz, for both R_1 and R_2. Only a subset of high-quality data is considered; measurements with singular values ≥ 0.75 for R_1 and ≥ 0.85 for R_2 are used from the method of Laske *et al.* (1994), and measurements with correlation coefficients $C_{zr} \geq 0.85$ for R_1 and $C_{zr} \geq 0.70$ for R_2 are used from DLOPy. The method of Laske *et al.* (1994) gives formal uncertainties for individual estimates, whereas DLOPy does not. Numbers in the top left denote correlation coefficients between the two datasets.

mometer that is routinely installed at GSN land stations. At this site, the water depth was about 4400 m, and the seismometer was emplaced in Ocean Drilling Project Hole 843B and 248 m below the seafloor. Site OSN1S had a Güralp CMG-3T broadband seismometer installed on the ocean floor, and OSN1B had the same sensor as OSN1S but pushed into the mud about 1 m deep. At the time this article was written, the metadata at the IRIS-DMC report an instrument orientation of 0° (i.e., alignment with the geographical coordinate system; or orientation unknown) for all three sensors.

The borehole installation at OSN1 was expected to have the lowest noise levels, particularly at frequencies below 10 mHz. Unfortunately, fluid circulation in the borehole caused noise levels to be unexpectedly high so that OSN1B actually produced the seismically quietest data (Stephen *et al.*, 2003). The borehole sensor outperformed the Güralp systems at frequencies beyond the microseism peak at frequencies above 0.3 Hz on the vertical components. The difference was even more striking for the horizontal components, for which the buried sensor at OSN1B clearly outperformed KS54000 at frequencies below 50 mHz, whereas KS54000 was clearly the best sensor at frequencies higher than 0.2 Hz. For our analysis, which is performed at low frequencies, we therefore expect to obtain the most stable results for OSN1B. We should also note

Table 2
A Comparison of Results for the Three Ocean-Bottom Seismometers
(OBSs) Deployed for the Ocean Seismic Network (OSN) Pilot
Experiment (Stephen *et al.*, 2003)

Method	OSN1 (°)	OSN1S (°)	OSN1B (°)
Laske *et al.* (1994)	126.41 ± 0.90	282.36 ± 2.24	147.99 ± 0.87
STACHPy	125.88 ± 1.87	281.28 ± 3.16	148.79 ± 2.73
DLOPy	126.72 ± 1.86	281.46 ± 3.64	148.57 ± 1.68
DLOPy, R_1 only	126.54 ± 1.86	281.03 ± 3.58	148.59 ± 1.72
DLOPy, R_2 only	127.91 ± 5.46	283.29 ± 11.52	148.42 ± 3.62

that burial greatly reduced the noise on the Güralp sensor at frequencies below 50 mHz, on all components, but had virtually no effect at high frequencies, though an improvement may be observed on the horizontal components at frequencies above 5 Hz.

Table 2 summarizes our results together with those obtained with the Laske *et al.* (1994) method and STACHPy, and the results are included in Figure 6. As expected, the error bars for the instrument orientation are the smallest for OSN1B and the largest for OSN1S using DLOPy and the Laske *et al.* (1994) method. It is unclear why the error is relatively large for OSN1B when using STACHPy, or why the error at OSN1S is larger for DLOPy than for STACHPy. The Laske *et al.* (1994) method produces the smallest error bars for each station. However, the orientations using the different methods place within the respective error bars so that discrepancies are statistically insignificant, for all three stations. We also present orientations obtained when using only R_1 or R_2 measurements. Because of the very small number of data, the error bars for the R_2 orientations are relatively large. Remarkably, the obtained instrument orientations still place within the error bars of the complete dataset. All these results indicate that, at least in the deep ocean, deployments of only several months are sufficient to orient the sensors reliably during postprocessing using long-period surface waves.

We also check the instrument orientation at the cabled ocean observatory MOBB (Romanowicz *et al.*, 2006). The cable allows GSN-like real-time access to the data. At this site, the water depth is 1000 m, and a Güralp CMG-1T was installed 0.5 m under the seafloor. The data can be obtained through the Northern California Earthquake Data Center (NCEDC), whereas some metadata are included at the IRIS-DMC in the metadata for the Berkeley Digital Seismograph Network (network code BK). An investigation of infragravity energy on the horizontal components of MOBB and OSN seismometers was the initial motivator for the work presented here (Doran and Laske, 2016) because understanding the directionality of infragravity waves requires precise knowledge of the instrument orientations.

We use data between 1 August 2011 and 1 September 2012, which compares to the duration of a typical modern temporary passive seismic OBS deployment. Using DLOPy, we obtain an orientation of 8.73° ± 1.82° using 575 measurements from 205 events, whereas STACHPy computes an

orientation of 7.36° ± 6.45° using 31 events. Choosing a slightly different time period from 1 November 2011 through 1 November 2012 gives 7.99° ± 1.78° from 564 measurements on 184 events in DLOPy and 6.92° ± 6.47° based on 28 events in the Stachnik *et al.* (2012) implementation. For both methods, the orientation varies on the order of 1°, which lies within the error bars and so is statistically insignificant. Using the Laske *et al.* (1994) method, we examined six frequencies between 8 and 23 mHz and obtained an orientation of 8.39° ± 1.11°. These results were calculated after excluding outliers during secondary screening. Higher frequencies were examined here than for ESK and SACV because of high noise at lower frequencies. Although in analyzing these data, we noticed that the Love-wave measurements produce more internally consistent results than the Rayleigh-wave measurements, with cleaner waveforms and lower uncertainties associated with the individual measurements. The Laske *et al.* (1994) method produces the smallest error bars, and STACHPy produces the largest, but all three results agree within their respective uncertainties. Discussion with MOBB station operators recently confirmed that a misorientation of 6°–8° is quite possible (B. Romanowicz, personal comm., 2016).

Compared to the results obtained for the OSN sensors, the error bars for STACHPy are surprisingly large, which we cannot fully explain at this point. DLOPy allows the inclusion of many more earthquakes, resulting in a drastic reduction of formal uncertainties. We do note that in raw-data plots similar to that in Figure 3, the fewer measurements cluster more tightly around an average value for OSN sensors than for the MOBB sensor for which we have vastly more data.

The actual orientation of the sensor is probably not known as the metadata at the IRIS-DMC and the NCEDC report an orientation of 0°. To conclude, at low-noise OBS sites, we should expect to be able to determine instrument orientations with error bars as small as 2° or less when using DLOPy. Whether this also means that we can determine instrument orientations to that accuracy is addressed in the final discussion at the end of this article.

OBSIP Deployments: The Hawaiian PLUME2 Deployment

We now examine a dataset obtained from one of the early OBSIP deployments: the Hawaiian PLUME2 (Laske
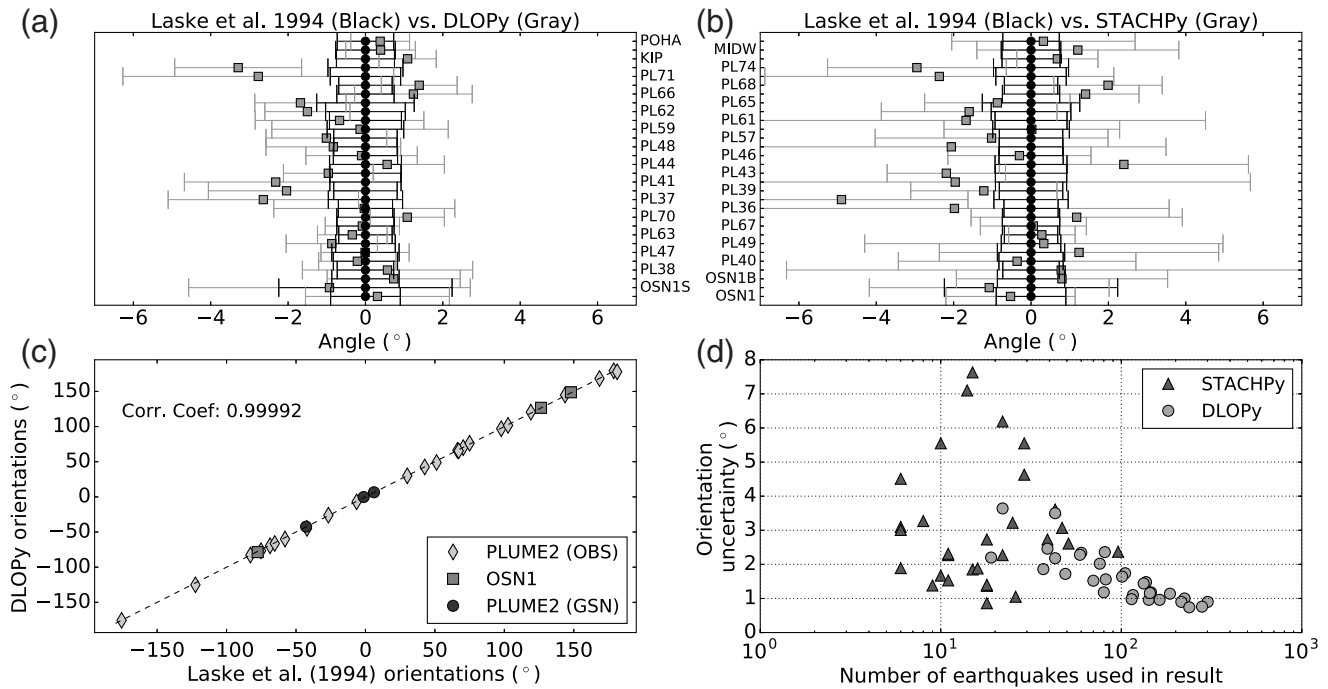
Figure 6. A comparison of results obtained with DLOPy, the method of Laske *et al.* (1994), and STACHPy for 30 stations of the Ocean Seismic Network (OSN) and Plume-Lithosphere Undersea Mantle Experiment (PLUME) deployments. (a,b) Comparison of the final instrument orientations with respect to those found using the Laske *et al.* (1994) method. (c) A comparison between Laske *et al.* (1994) and DLOPy. Phase 2 deployment for PLUME (PLUME2) GSN denotes GSN stations IU.KIP, IU.MIDW, and IU.POHA. (d) Comparison of the instrument orientation uncertainties obtained using DLOPy to those of STACHPy.

et al., 2009; Rychert *et al.*, 2013). The PLUME2 deployment provided useful three-component seismograms for 24 broadband OBSs and 10 land stations from May 2006 through June 2007. Here we concentrate on the OBSs. The sensors were Güralp CMG-3T sensors in WHOI OBS packages and Nanometrics Trillium T-240s in the SIO packages. All OBSs were deployed at water depths between 4500 and 5800 m on and off the Hawaiian Swell. The temporary PLUME deployment was complemented by three permanent GSN stations: Kipapa (IU.KIP; an STS-1 sensor), Pohakuloa (IU.POHA; a KS54000), and Midway Island (IU.MIDW; an STS-2). For these 27 stations, we use our new technique to determine the instrument orientation and compare our results with those of Rychert *et al.* (2013), which used the Laske *et al.* (1994) high-quality benchmark method (published in their online supplement).

Figure 6 displays a comprehensive summary of the results. There is excellent agreement between the results using DLOPy and the Laske *et al.* (1994) method (see Fig. 6a,c). The latter yields consistently smaller error bars than DLOPy but, except for station PL74, all of our results agree with those obtained with the Laske *et al.* (1994) method, to within our error bars but not necessarily the Laske error bars. Station PL74 shows a discrepancy of about 4°. The sensor is nearly naturally polarized so that the BH1 component is almost aligned with geographic north. Upon inspection of Laske's original arrival-angle data, we notice that the six Rayleigh-wave frequencies yield an orientation result similar to what we get with the new

method (−2°), but the Love-wave frequencies consistently have the opposite sign, are internally more consistent and therefore dominate the averaging process, leading to an orientation of about +2°. Because STACHPy also only uses Rayleigh waves, this method yields a similar inconsistency for PL74 against the Laske *et al.* (1994) method but agrees with the results from DLOPy. A discrepancy of 4° seems unacceptable, but note that this is an exception in a set of 27 stations. We note that seven SIO OBS orientations published by Rychert *et al.* (2013) differ from our results by ±180°. The method of Laske *et al.* (1994) currently does not distinguish between δ and δ ± 180°. A comparison with synthetic seismogram could remove this ambiguity but is not implemented. The stations affected are PL38, PL40, PL47, PL49, PL63, PL67, and PL70 (for more information, see Ⓔ Table S2, available in the electronic supplement to this article).

Overall, STACHPy yields much larger error bars, sometimes well over 5°, and the instrument orientations seem to agree less well with those by Laske. Nevertheless, the differences are never statistically significant. It is also revealing to investigate the dependency of the orientation error bars on the number of unique earthquakes used. Our new method consistently uses many more earthquakes than STACHPy, by an order of magnitude. As discussed earlier, measuring multiple frequencies and multiple wavetrains allows a much greater number of events to be utilized. This usually yields significantly smaller error bars, by a factor of nearly 2 on average.

Table 3
Summary Comparison for the Cascadia Initiative OBSs
between Results of DLOPy, STACHPy, and Those
Disseminated by the OBS Instrument Pool (OBSIP)

Method	Deployments	Number of Stations	Median Error (°)	Median Number of Events
DLOPy	1–4	223	4.92	31
OBSIP	1–3	161	17.0	10
STACHPy	1–4	213	5.73	7

Results for the OBSs of the Cascadia Initiative

We now turn our attention to the four 1-year OBS deployments of the CI community experiment (Toomey *et al.*, 2014). CI was an amphibious experiment, that is, there was a land component, but we analyze only the OBS records using both DLOPy as well as STACHPy. The complete table of instrument orientations for 223 OBSs can be found in Ⓔ Table S1, whereas Table 3 provides an overall summary. As of 23 May 2016, the IRIS-DMC MetaData Aggregator contained no information on the azimuths of CI horizontal channels. Lodewyk and Sumy (2015) report that all data have been converted to a left-hand coordinate system, with H2 being 90° clockwise from H1, as of 2 May 2014. We accessed the data between 30 October 2015 and 28 March 2016 and assume this convention.

Applying DLOPy to OBS records of all four deployments, we obtain instrument orientations and related uncertainties for 223 stations. With STACHPy, we obtain orientations for only 213 stations. At the time this article was written, IRIS OMO disseminates reports containing estimated instrument orientations for the first three deployments (Lodewyk and Woodward, 2014; Lodewyk *et al.*, 2014; Lodewyk and Sumy, 2015). Although their number of events per station generally agrees with that using STACHPy, OMO obtains much larger error bars than we get with both Python codes. We speculate that this is a result of OMO's usage of the standard deviation of the data as error bars instead of determining a smaller error of the mean (Sumy *et al.*, 2015). As documented in Figure 7, most of our errors are about 6° or smaller, with a median of 4.92° (Table 3). Our minimum error is 1.68°.

Having a closer look at individual stations, a wide range of errors in the station orientation may be due to widely different ambient noise levels at the OBSs as a result of deployment in a wide range of water depths between less than 100 m and nearly 4500 m. As found by many other investigators of OBS data, shallow-water environments tend to be significantly noisier than those in the deep ocean, particularly at long periods. Consequently, our error bars at shallow stations (water depth $z \leq 500$ m) are three or more times greater than those at deeper stations (Fig. 8). This tends to be the case for all three algorithms. There also appears to be a relationship between the number of events useful for analysis and water depth, for all three implementations, supporting the idea that deep ocean environments are quieter at long
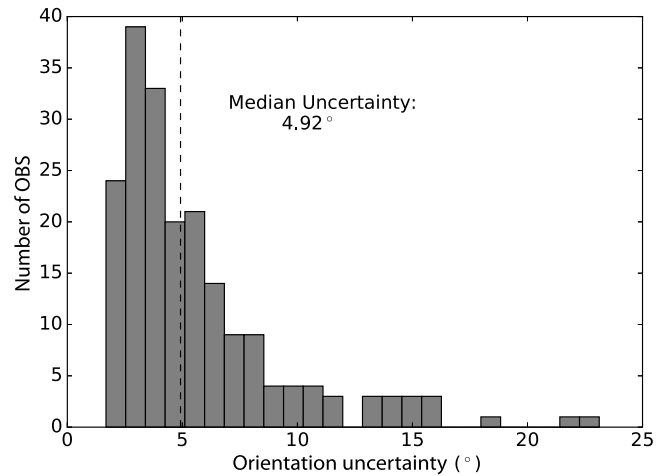

Figure 7. Histogram of station orientation uncertainties obtained for the four 1-year CI deployments using DLOPy. Most uncertainties are smaller than 7°, though a few uncertainties in our database are larger than 100° (not shown).

periods (25–100 s, as used in this study). As already documented in Table 3 and also obvious from Figure 8, our results are based on many more earthquakes than those using other algorithms. DLOPy is most consistent in assigning larger error bars for results that are based on only a few events. Recall that the two other implementations are based on purely statistical treatment of data samples, which can result in unrealistically small error bars for small datasets, unlike our bootstrapping approach.

We also investigate the dependence of the orientation error on the OBS design. For the CI, many instruments were newly designed to serve two purposes. First, for deployment on the continental shelf, the instruments needed to be trawl resistant. The danger of OBSs being disturbed or even displaced by fishing trawlers is particularly great on the wider shelf along the U.S. east coast, but also exists along the narrower west coast shelf. Second, the OSNPE clearly showed that burial of a sensor reduces long-period noise in the seismic records. Burial of the sensors is prohibitively costly and time consuming and was not an option for the CI, but some type of shielding of the sensor should go a long way to reduce current- and wave-induced noise, particularly at shallow sites. Although all three OBSIP IICs provided instruments, LDEO designed the trawl resistant mount and SIO designed the Abalone. All three groups used the Nanometrics Trillium Compact 120 as sensor. In addition, WHOI provided traditional broadband OBSs using a Güralp CMG-3T sensor that were deployed primarily in the deep ocean at water depths greater than 2000 m. Furthermore, two traditional SIO instruments were deployed in year 4 (see Data and Resources for additional information). Figure 9 shows orientation uncertainties as a function of depth and instrument type. As already mentioned, errors tend to be larger for shallow deployments, particularly at water depths less than 200 m. Some large-error deeper deployments tend to be ones
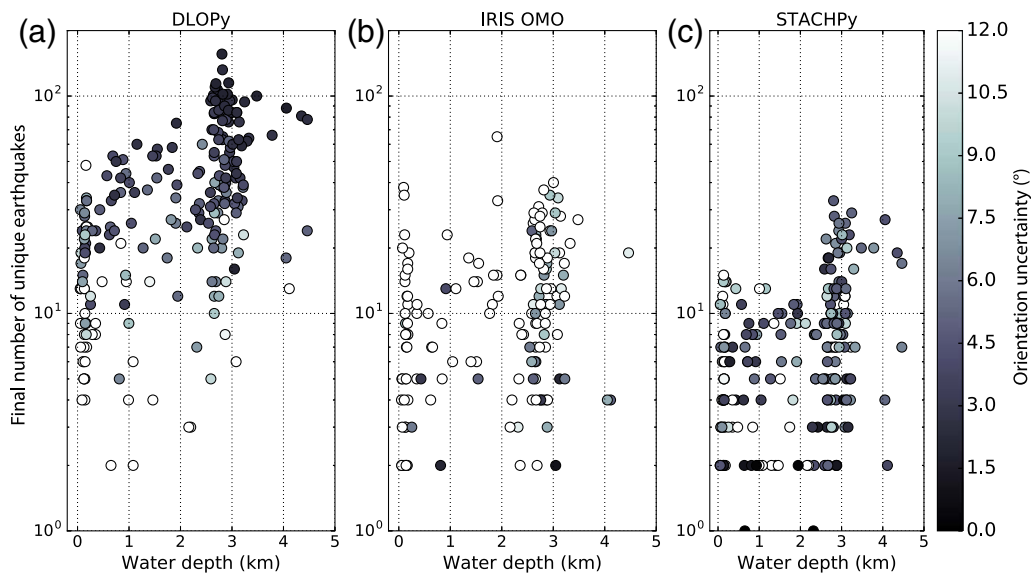
Figure 8.   Number of unique earthquakes utilized (*y* axis) and uncertainty of CI station orientations (gray shades) as a function of water depth. (a) Results using DLOPy; (b) results reported by IRIS OMO; and (c) results using STACHPy. The color version of this figure is available only in the electronic edition.

with technical problems and/or premature shutdown. For example, for WHOI's year-4 station FS16D deployed at 1080-m water depth, we have only three useful measurements from two events. Although hoping that shielding the seismic sensor would help reduce noise levels, particularly on the horizontal components in shallow water, we cannot document that, because no standard OBS for comparison was deployed at depths shallower than 716 m. At this depth, the WHOI year-4 instrument at G09D (for which we have 63 data from 24 earthquakes) performed just as well as the shielded OBSs. However, remarkably and perhaps somewhat unexpectedly, the SIO Abalone has nearly consistently smaller orientation errors than most other instruments for larger water depths. A relatively large error of 15° at year-2 SIO site G12B is the result of a dearth of useful data, with only 17 measurements from six earthquakes. Although our error bars are not a direct measure for ocean noise on the horizontal seismmometer components, we are confident that we can use it as a proxy. It seems therefore that the SIO Abalone competes well with more traditional OBS packages in the deep ocean along Cascadia. Our results are consistent with those of Bell *et al.* (2015), who examine tilting on the horizontal components of OBSs from Cascadia year-1 deployment at long periods (50 s) and conclude that SIO shielded instruments are quieter than unshielded OBSs at similar depths.

As with PLUME, we conduct a preliminary investigation of the internal consistency of our CI measurements. Figure 10 shows 30-mHz R_1 arrival-angle measurements for two events in the Solomon Islands that occurred within some 20 km of each other during the year-3 deployment. These events produced high-quality ($C_{zr} \geq 0.90$) measurements on nearly every station. The progression of the arrival-angle anomaly across the CI network is spatially coherent

and gives consistent measurements between the events, except for a small number of outliers. Arrivals are systematically rotated clockwise in the northwest of the network but counterclockwise in most of the rest, documenting complex lateral refraction of the Rayleigh waves between the source region and the network. Wavefield complexity as expressed by surface-wave arrival-angle anomalies has recently been documented on the larger USArray Transportable Array (Foster, Ekström, and Hjörleifsdóttir, 2014).

Comparison with Other Seismic Techniques

Several other seismic approaches have been utilized to determine the station orientation during data postprocessing. A number of studies analyze *P*-wave particle motion to estimate horizontal orientations (e.g., Toomey *et al.*, 1998; Schulte-Pelkum *et al.*, 2001; Niu and Li, 2011; Stachnik *et al.*, 2012), either alone or in combination with surface-wave arrival angles. Although both approaches should and often do produce consistent results, we prefer to analyze surface-wave arrival angles for a number of reasons. First and foremost, *P*-wave arrivals tend to have a small amplitude and are much more likely to be contaminated by noise than surface-wave arrivals, particularly on OBS records. Stachnik *et al.* (2012) found a factor of 5 more events with acceptable SNRs using surface waves as opposed to *P*-waves. Another reason against teleseismic body-wave arrivals is that events have to be in a specific distance range to avoid interference with other phases. This dramatically narrows the number of suitable events compared to a surface-wave study, thus limiting much-needed azimuthal coverage. Not the last, long-period *P* waves can display 10° or more deviation from the source back azimuth (Schulte-Pelkum *et al.*, 2001), and thus,
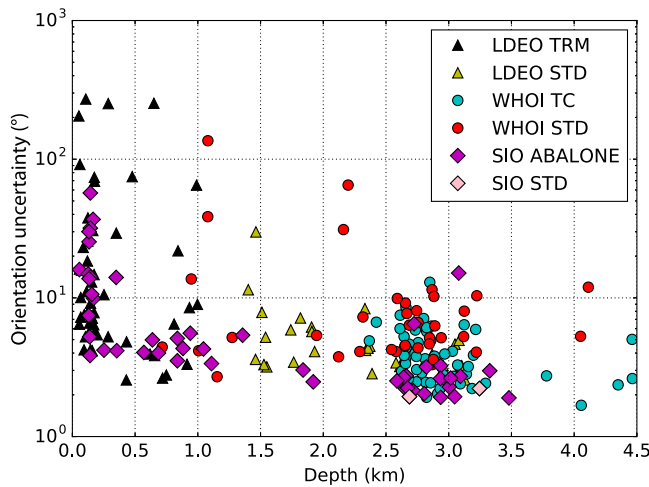
Figure 9. Instrument orientation uncertainties determined using DLOPy for the CI OBSs, as a function of water depth and sorted by instrument type. Scripps Institution of Oceanography (SIO) Abalone and Lamont–Doherty Earth Observatory (LDEO) trawl-resistant-mount (TRM) are shielded, trawl-resistant instruments. LDEO TRM OBSs were deployed only in shallow water ($z < 1000$ m). LDEO STDs were standard LDEO OBSs. Woods Hole Oceanographic Institute (WHOI) Trillium Compacts (TCs) were newly designed for CI deployments. All these used a Nanometrics TC 120s as sensor. The WHOI STD OBSs used a broadband Güralp CMG-3T sensor, and the SIO STD OBSs used a Nanometrics Trillium 240 sensor. The color version of this figure is available only in the electronic edition.

as waves propagating in the 3D Earth may accumulate more effects from lateral refraction than surface waves propagating on a 2D surface.

Other seismic techniques that determine the instrument orientation during postprocessing include shear-wave splitting and RF analyses. Janiszewski and Abers (2015) utilize the RF method, whereby the seismometer orientation is determined by minimizing the RF energy on the transverse component, to orient a number of Cascadia Deployment 1 and 3 stations. Zha et al. (2013) use ambient noise correlations to orient the instruments of an OBS network in the Eastern Lau Spreading Center (Zha et al., 2014). This method calculates the three-component empirical Green's function from ambient seismic noise and determines the arrival angle based on principles similar to those described in our article. Regardless of the approach, all methods should yield similar values for the instrument orientation to within their respective error bars. Figure 11 compares the results from these methods with those obtained with DLOPy. It should be mentioned here that Zha et al. (2013) use a right-handed coordinate system in their representation and we converted their published results by subtracting them from 180°. Although there is general agreement in the instrument orientations between using these methods and DLOPy, disagreements are much larger than those between using DLOPy and the surface-wave arrival-angle method of Laske et al. (1994; see Fig. 6). Some of the discrepancy may stem from the fact

that these methods are based on different seismic observables, and therefore are subject to different errors and biases. However, all methods should give statistically compatible results because there is only one true orientation of the instrument (assuming it does not vary with time), and the methods discussed here show some significant disagreement. Perhaps these examples support the idea that at least for OBSIP deployments, it may be beneficial that instrument orientations are determined using consistently one and the same technique.

Discussion

This article introduces an automated technique to measure Rayleigh-wave arrival angles on OBS deployments, with the ultimate goal of determining the orientation of the horizontal seismometer components. Obtaining instrument orientations through the analysis of surface waves has several advantages over using other approaches: surfaces waves are typically the largest signal in a seismogram, so SNRs are typically best. There are fewer restrictions on earthquake choices compared to other teleseismic techniques, usually leading to a better azimuthal event coverage. This is important in light of the fact that the arrival direction of seismic phases is typically influenced by lateral refraction in the heterogeneous Earth. For stations with only a few earthquakes, particularly when they cluster in a certain source region, the resulting instrument orientation may be in error by as much as 5° or even larger. To obtain unbiased instrument orientations, the arrival angles should therefore undergo a nonlinear joint inversion for structure and instrument orientation (e.g., Laske, 1995; Larson and Ekström, 2002). Here, we did not do this but simply determined the medians in our datasets. DLOPy tries to accommodate this omission through larger error bars that result automatically for smaller datasets when applying the bootstrap process.

Larson and Ekström (2002) determined surface-wave arrival angles for GSN stations and other permanent, observatory-quality seismic stations using events between 1989 and 1998 and then compared the differences in the resulting station orientations when including or omitting Earth structure in a joint inversion. For stations with more than 100 measurements, the median discrepancy between both results is only 0.2°. But when the number of measurements drops below 100, the discrepancy can be much larger, and the median in the discrepancy for the stations in Larson and Ekström (2002) that fall in this category increases to 1.1°. For GSN stations, and using the 10-year dataset, the number of high-quality measurements can drop to below 10 but the discrepancy between simple average and result from joint inversion never reaches beyond 5.6°. With 23 high-quality data, station SMTC (Superstition Mountain of the Southern Californian TERRAscope/Trinet Network) had the largest discrepancy at 5.6°. However, the error obtained from the inversion was 6.5° so this discrepancy is statistically insignificant. We also demonstrated that the repeatability of indi-
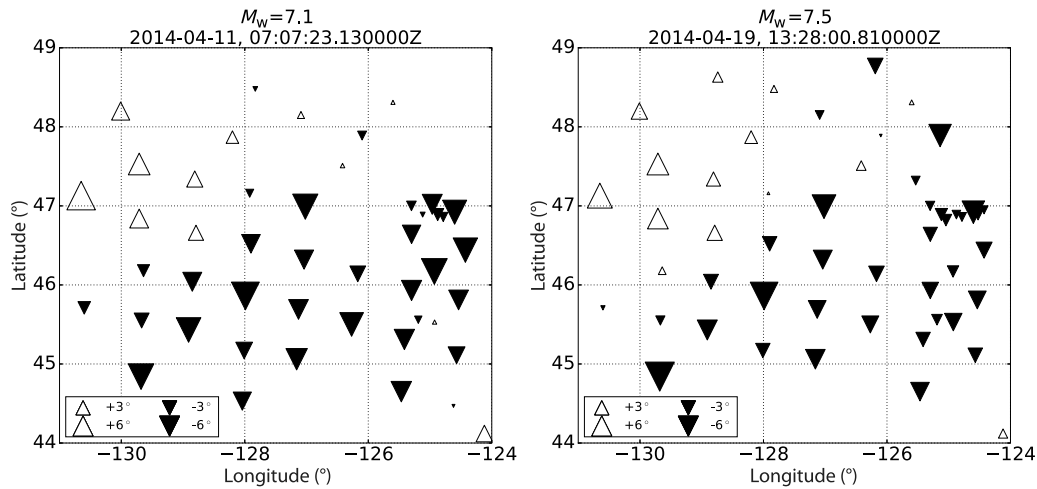
Figure 10. Arrival-angle deviations from the source–receiver great circle for R_1 measured at 30 mHz for two close events during the year-3 CI deployment. The events occurred in April 2014 in the Solomon Islands region, at epicentral distances of about 85° (from center of network), with a back azimuth of about −105°. The events were large, with M_w 7.1 and 7.5, and occurred about eight days apart. Only stations with a high cross-correlation value (CC > 0.9) are shown. The individual measurements were corrected for the final instrument orientation. The shapes of the symbols represent arrival-angle deviations in degrees from the great-circle path. Clockwise rotations are shown in white, whereas counterclockwise rotations are shown in black. The triangles are scaled by the magnitude of the deviations.

vidual measurements using the same or different techniques can result in differences in instrument orientations on the order of 1°–2°. The main point here is that, for temporary OBS deployments and studies that do not perform a joint inversion for structure to determine instrument orientations, we should probably expect a possible bias on the order of a few degrees. We therefore feel that increasing the error bar from 2-σ (as reported in other studies) to 4-σ (this study) is well justified.

Both Laske (1995) and Larson and Ekström (2002) analyze both Rayleigh and Love waves. Because Love waves sense shallower structure at the same frequency as Rayleigh waves (e.g., Laske and Widmer-Schnidrig, 2015), some argue that Love waves in general are subject to greater lateral refraction (e.g., G. Masters, personal comm., 2016). Others argue that Love waves are subject to severe interference with Rayleigh-wave overtones (e.g., Foster, Nettles, and Ekström, 2014; Jin and Gaherty, 2015) and should therefore be excluded from analysis. We find that high singular values in the Laske *et al.* (1994) method, together with the ellipticity of the particle motion, are excellent indicators for the purity of Love waves, and we maintain that including Love waves in the process yields less bias in station orientations. In fact, it is not uncommon that the high-quality Love-wave arrival-angle datasets are larger than the corresponding Rayleigh-wave datasets. In a noisy seismogram, it is often easier to make high-quality Love-wave measurements than Rayleigh-wave measurements using the Laske *et al.* (1994) method because the waveforms are less dispersed in the frequency band of interest. Consequently, windows are shorter, and less noise or other signals enter the measurement. We attempted to analyze Love waves with DLOPy. Using the dispersion maps of Ma *et al.* (2014) to determine our analysis windows, we would expect that the correct angle is found when the rotated H2 component H2′ has no correlation with the Hilbert-trans-

formed vertical (i.e., it is zero). It turns out that in this case, we again achieve maximum correlation between H1′ and the Hilbert-transformed vertical, so we used the latter as criterion. However, Rayleigh-wave overtone signals in the same window also lead to increased correlations between these components. To make matters worse, random checks of the corresponding three-component seismograms indicate that in highly correlating cases, the front end of the Rayleigh wave is present, implying that our automated windows are too long to separate out the Love wave effectively. We therefore exclude the analysis of Love waves at this point.

The exclusion of Love waves is an important weakness of Rayleigh-wave only methods such as DLOPy and may explain much of the discrepancy seen in the final orientation estimate at ESK. Also, if we deal with uneven event coverage on a heterogeneous Earth, then orientations based solely on statistical averages instead of on a joint inversion for structure may well be biased. Considering only the year 2004 and only statistical averages, the Laske *et al.* (1994) method calculates an overall orientation of $0.17° \pm 0.82°$ for ESK, but using only Rayleigh waves yields an orientation of $1.32° \pm 1.97°$, which is statistically consistent with the DLOPy results. Although ESK may be an extreme case, a systematic error of some 2° must be considered a realistic possibility. Using only Rayleigh-wave data is probably a sufficient and appropriate strategy for quick determination of instrument orientations for OBS and temporary networks. Permanent observatory stations such as those of the GSN should include high-quality Love-wave data to reduce potential bias.

Summary

We developed an automated method DLOPy to determine the geographic orientation of horizontal seismometer
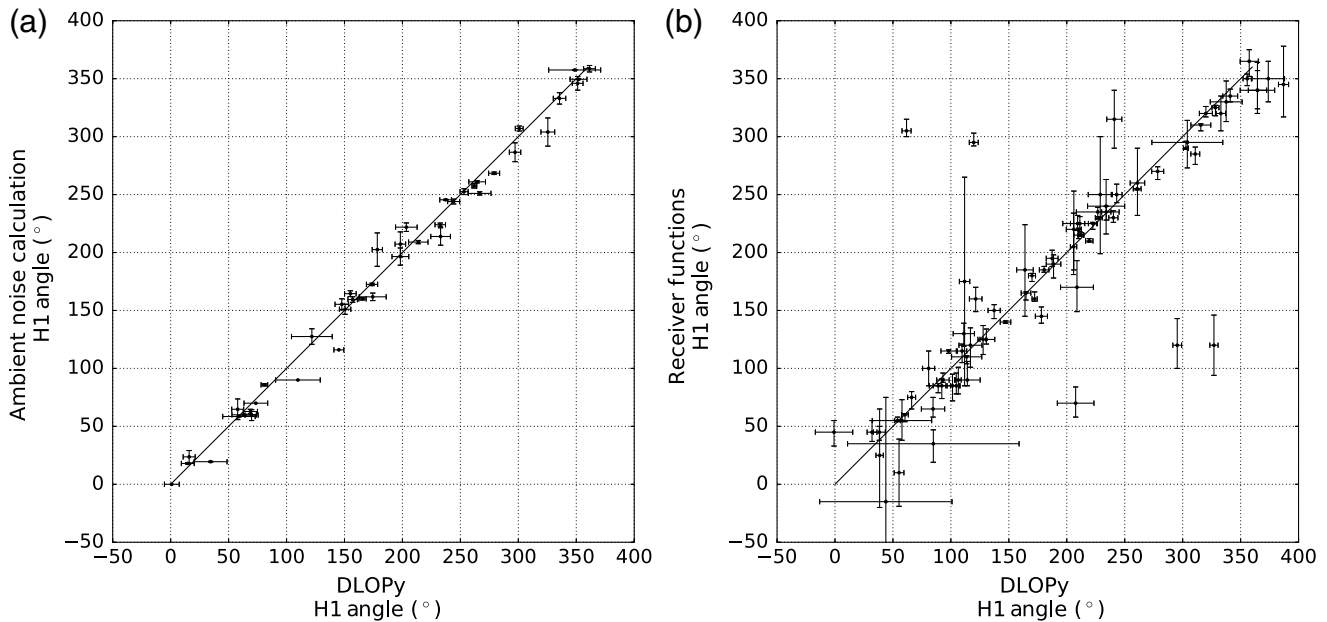
Figure 11.   A comparison of results obtained for OBS deployments using DLOPy and two other methods. (a) Instrument orientations for the Eastern Lau Spreading Center deployment using ambient noise correlations (Zha *et al.*, 2013, 2014). (b) Instrument orientations for the CI year-1 and year-3 deployments using receiver functions (Janiszewski and Abers, 2015). Both case comparisons exhibit significant disagreement. Compare these results with those shown in Figure 6c.

components. Our automated procedure measures frequency-dependent arrival angles of intermediate- and long-period fundamental-mode Rayleigh waves from the first major and minor great-circle arcs. Our technique uses optimized but conservative confidence intervals to provide realistic orientation estimates and uncertainties and can produce reliable sensor orientations for deployments as short as several months. We validated our method by comparing our overall instrument orientations as well as our individual arrival-angle measurements with results from well-established techniques at high-quality GSN stations and OBS deployments of various lengths. Although DLOPy can be applied to all broadband three-component instruments, it is particularly intended for use with free-fall OBSs and other temporary deployments in need of quick determination of instrument orientations. We presented orientations for all OBS stations with sufficient data from the large-scale CI amphibious project. We provided evidence that newly developed shielded OBSs provide substantial improvement to long-period data quality even in the deep ocean. Our technique was designed with accuracy, automation, and ease-of-use in mind. DLOPy provides a simple, stable, and standardized method to compute accurate orientations and meaningful uncertainties with minimal bias without undergoing a formal joint inversion for Earth structure.

Data and Resources

Waveforms, station metadata, and earthquake catalogs are publicly available and were accessed through the Incorporated Research Institutions for Seismology Data Manage-

ment Center (IRIS-DMC; http://ds.iris.edu/ds/nodes/dmc/, last accessed May 2016) and the Northern California Earthquake Data Center (NCEDC; doi: 10.7932/NCEDC) for station MOnterey Bay Broadband Observatory (MOBB; http://www.ncedc.org/, last accessed May 2016). Global Seismographic Network (GSN) operator IRIS-IDA provided the data for stations ESK (Eskdalemuir, Scotland) and SACV (Santiago Island, Cape Verde). IRIS Ocean-Bottom Seismometer (OBS) Management Office (OMO) provides the MATLAB (www.mathworks.com/products/matlab, last accessed October 2016) implementation of the Stachnik *et al.* (2012) code on the OBS Instrument Pool (OBSIP) website at http://www.obsip.org/data/obs-horizontal-orientation/ (last accessed May 2016). More information about the Cascadia Initiative (CI) can be found at http://cascadia.uoregon.edu (last accessed May 2016). More information on OBSIP and CI instrumentation can be found at http://www.obsip.org/instruments/cascadia-instruments/ (last accessed September 2016). Zhitu Ma provides his surface-wave dispersion maps on the LITHO1.0 website at http://igppweb.ucsd.edu/~gabi/litho1.0.html (last accessed April 2016). Many of the geodesic calculations utilize the methods of Karney (2013), and implementations in several languages can be found at http://geographiclib.sourceforge.net (last accessed March 2016). Figure 1 was drawn using Adobe Illustrator, whereas all other figures were drawn using Python 2.7 (Python Software Foundation, http://www.python.org, last accessed September 2016). A distribution of our code will be available at http://igppweb.ucsd.edu/~gabi/ (last accessed October 2016) and http://igppweb.ucsd.edu/~adoran/ (last accessed October 2016). The full table of Cascadia orientations and a table

of phase 2 deployment for Plume-Lithosphere Undersea Mantle Experiment (PLUME2) orientations can be found in Ⓔ Tables S1 and S2, available in the electronic supplement to this article, as well as the websites above.

Acknowledgments

References

Ahern, T., R. Casey, D. Barnes, R. Benson, and T. Knight (2012). *SEED Reference Manual, Version 2.4 IRIS,* https://www.fdsn.org/seed_manual/SEEDManual_V2.4.pdf (last accessed September 2016).

Anderson, P., F. Duennebier, and R. Cessaro (1987). Ocean borehole horizontal seismic sensor orientation determined from explosive charges, *J. Geophys. Res.* **92,** no. B5, 3573–3579, doi: 10.1029/JB092iB05p03573.

Baker, G. E., and J. Stevens (2004). Backazimuth estimation reliability using surface wave polarization, *Geophys. Res. Lett.* **31,** L09611, doi: 10.1029/2004GL019510.

Bell, S. W., D. W. Forsyth, and Y. Ruan (2015). Removing noise from the vertical component records of ocean-bottom seismometers: Results from year one of the Cascadia Initiative, *Bull. Seismol. Soc. Am.* **105,** no. 1, 300–313, doi: 10.1785/0120140054.

Beyreuther, M., R. Barsch, L. Krischer, T. Megies, Y. Behr, and J. Wassermann (2010). ObsPy: A Python toolbox for seismology, *Seismol. Res. Lett.* **81,** no. 3, 530–533, doi: 10.1785/gssrl.81.3.530.

Chael, E. (1997). An automated Rayleigh-wave detection algorithm, *Bull. Seismol. Soc. Am.* **87,** no. 1, 157–163.

Collins, J., F. Vernon, J. Orcutt, R. Stephen, K. Peal, F. Wooding, F. Spiess, and J. Hildebrand (2001). Broadband seismology in the oceans: Lessons from the Oceans Seismic Network pilot experiment, *Geophys. Res. Lett.* **28,** no. 1, 49–52.

Doran, A., and G. Laske (2016). Infragravity waves and horizontal seafloor compliance, *J. Geophys. Res.* **121,** no. 1, 260–278, doi: 10.1002/2015JB012511.

Duennebier, F., P. Anderson, and G. Fryer (1987). Azimuth determination of and from horizontal ocean bottom seismic sensors, *J. Geophys. Res.* **92,** no. B5, 3567–3572, doi: 10.1029/JB092iB05p03567.

Dziewonski, A., and D. Anderson (1981). Preliminary reference earth model, *Phys. Earth Planet. In.* **25,** no. 4, 297–356.

Efron, B. (1979). Bootstrap methods: Another look at the jackknife, *Ann. Stat.* **7,** no. 1, 1–26.

Ekström, G., and R. W. Busby (2008). Measurements of seismometer orientation at USArray Transportable Array and backbone stations, *Seismol. Res. Lett.* **79,** no. 4, 554–561, doi: 10.1785/gssrl.79.4.554.

Foster, A., G. Ekström, and V. Hjörleifsdóttir (2014). Arrival-angle anomalies across the USArray Transportable Array, *Earth Planet. Sci. Lett.* **402,** 58–68, doi: 10.1016/j.epsl.2013.12.046.

Foster, A., M. Nettles, and G. Ekström (2014). Overtone interference in array-based Love-wave phase measurements, *Bull. Seismol. Soc. Am.* **104,** no. 5, 2266–2277, doi: 10.1785/0120140100.

Harris, F. (1978). On the use of windows for harmonic analysis with the discrete Fourier transform, *Proc. IEEE* **66,** no. 1, 51–83, doi: 10.1109/PROC.1978.10837.

Janiszewski, H., and G. Abers (2015). Imaging the plate interface in the Cascadia seismogenic zone: New constraints from offshore receiver functions, *Seismol. Res. Lett.* **86,** no. 5, 1261–1269, doi: 10.1785/0220150104.

Jin, G., and J. Gaherty (2015). Surface wave phase-velocity tomography based on multichannel cross-correlation, *Geophys. J. Int.* **201,** no. 3, 1383–1398, doi: 10.1093/gji/ggv079.

Karney, C. F. F. (2013). Algorithms for geodesics, *J. Geodes.* **87,** no. 1, 43–55, doi: 10.1007/s00190-012-0578-z.

Krischer, L., T. Megies, R. Barsch, M. Beyreuther, T. Lecocq, C. Caudron, and J. Wassermann (2015). ObsPy: A bridge for seismology into the scientific Python ecosystem, *Comput. Sci. Discov.* **8,** no. 1, 014003, doi: 10.1088/1749-4699/8/1/014003.

Larson, E., and G. Ekström (2002). Determining surface wave arrival angle anomalies, *J. Geophys. Res.* **107,** no. B6, doi: 10.1029/2000JB000048.

Larson, E., J. Tromp, and G. Ekström (1998). Effects of slight anisotropy on surface waves, *Geophys. J. Int.* **132,** 654–666.

Laske, G. (1995). Global observation of off-great-circle propagation of long-period surface waves, *Geophys. J. Int.* **123,** no. 1, 245–259, doi: 10.1111/j.1365-246X.1995.tb06673.x.

Laske, G., and G. Masters (1996). Constraints on global phase velocity maps from long-period polarization data, *J. Geophys. Res.* **101,** no. B7, 16,059–16,075, doi: 10.1029/96JB00526.

Laske, G., and G. Masters (1998). Surface-wave polarization data and global anisotropic structure, *Geophys. J. Int.* **132,** no. 3, 508–520, doi: 10.1046/j.1365-246X.1998.00450.x.

Laske, G., and R. Widmer-Schnidrig (2015). Theory and observations: Normal mode and surface wave observations, in *Treatise on Geophysics,* Second Ed., G. Schubert (Editor-in-Chief), Vol 1, Elsevier, Oxford, United Kingdom, 117–167.

Laske, G., J. Collins, C. Wolfe, S. Solomon, R. Detrick, J. Orcutt, D. Berovici, and E. Hauri (2009). Probing the Hawaiian hot spot with new ocean bottom instruments, *Eos Trans. AGU* **90,** no. 41, 362–363, doi: 10.1029/2009EO410002.

Laske, G., G. Masters, and W. Zurn (1994). Frequency-dependent polarization measurements of long-period surface waves and their implications for global phase-velocity maps, *Phys. Earth Planet. In.* **84,** 111–137.

Lodewyk, J., and D. Sumy (2015). Cascadia amphibious array ocean bottom seismograph horizontal orientations, *2013-2014 OBS Deployments, Version 2.1,* OBSIP Management Office, IRIS, 24 pp.

Lodewyk, J., and B. Woodward (2014). Cascadia amphibious array ocean bottom seismograph horizontal orientations, *2012-2013 OBS Deployments, Version 1.0,* OBSIP Management Office, IRIS, 25 pp.

Lodewyk, J., A. Frassetto, A. Adinolfi, and B. Woodward (2014). Cascadia amphibious array ocean bottom seismograph horizontal orientations, *2011-2012 OBS Deployments, Version 3.0,* OBSIP Management Office, IRIS, 23 pp.

Ma, Z., and G. Masters (2014). A new global Rayleigh- and Love-wave group velocity dataset for constraining lithosphere properties, *Bull. Seismol. Soc. Am.* **104,** no. 4, 2007–2026, doi: 10.1785/0120130320.

Ma, Z., G. Masters, G. Laske, and M. Pasyanos (2014). A comprehensive dispersion model of surface wave phase and group velocity for the globe, *Geophys. J. Int.* **199,** no. 1, 113–135, doi: 10.1093/gji/ggu246.

Megies, T., M. Beyreuther, R. Barsch, L. Krischer, and J. Wassermann (2011). ObsPy—What can it do for data centers and observatories? *Ann. Geophys.* **54,** no. 1, 47–58, doi: 10.4401/ag-4838.

Niu, F., and L. Li (2011). Component azimuths of the CEArray stations estimated from *P*-wave particle motion, *Earthq. Sci.* **24,** no. 1, 3–13, doi: 10.1007/s11589-011-0764-8.

Riedel, M., M. Cote, P. Neelands, G. Middleton, G. Standen, R. Iuliucci, M. Ulmi, C. Stacey, R. Murphy, D. Manning, *et al.* (2014). 2012 Haida

Gwaii M_w 7.7 earthquake response—Ocean bottom seismometer re-location and geophone orientation analysis and quality control, *Geol. Surv. of Canada Open-File 7632*, 79 pp., doi: 10.4095/295551.

Ringler, A. T., C. R. Hutt, K Persefield, and L. S. Gee (2013). Seismic station installation orientation errors at ANSS and IRIS/USGS Stations, *Seismol. Res. Lett.* **84,** no. 6, 926–931, doi: 10.1785/0220130072.

Romanowicz, B., D. Stakes, D. Dolenc, D. Neuhauser, P. McGill, R. Uhrhammer, and T. Ramirez (2006). The Monterey Bay broadband ocean bottom seismic observatory, *Ann. Geophys.* **49,** nos. 2/3, 607–623, doi: 10.4401/ag-3132.

Rychert, C., G. Laske, N. Harmon, and P. Shearer (2013). Seismic imaging of melt in a displaced Hawaiian plume, *Nature Geosci.* **6,** 657–660, doi: 10.1038/ngeo1878.

Schulte-Pelkum, V., G. Masters, and P. Shearer (2001). Upper mantle anisotropy from long-period *P* polarization, *J. Geophys. Res.* **106,** no. B10, 21,917–21,934, doi: 10.1029/2001JB000346.

Selby, N. D. (2001). Association of Rayleigh waves using backazimuth measurements: Application to test ban treaty verification, *Bull. Seismol. Soc. Am.* **91,** no. 3, 580–593, doi: 10.1785/0120000068.

Stachnik, J. C., A. F. Sheehan, D. W. Zietlow, Z. Yang, J. Collins, and A. Ferris (2012). Determination of New Zealand ocean bottom seismometer orientation via Rayleigh-wave polarization, *Seismol. Res. Lett.* **83,** no. 4, 704–712, doi: 10.1785/0220110128.

Stephen, R., F. Spiess, J. Collins, J. Hildebrand, J. Orcutt, K. Peal, F. Vernon, and F. Wooding (2003). Ocean Seismic Network pilot experiment, *Geochem. Geophys. Geosys.* **4,** no. 10, 1–38, doi: 10.1029/2002GC000485.

Sumy, D. F., J. A. Lodewyk, R. L. Woodward, and B. Evers (2015). Ocean-bottom seismograph performance during the Cascadia Initiative, *Seismol. Res. Lett.* **86,** 1238–1246, doi: 10.1785/0220150110.

Toomey, D., R. M. Allen, A. H. Barclay, S. W. Bell, P. D. Bromirski, R. L. Carlson, X. Chen, J. A. Collins, R. P. Dziak, B. Evers, *et al.* (2014). The Cascadia Initiative: A sea change in seismological studies of subduction zones, *Oceanography* **27,** no. 2, 138–150, doi: 10.5670/oceanog.2014.49.

Toomey, D., W. Wilcock, S. Solomon, W. Hammond, and J. Orcutt (1998). Mantle seismic structure beneath the melt region of the East Pacific Rise from *P* and *S* wave tomography, *Science* **280,** 1224–1227, doi: 10.1126/science.280.5367.1224.

Webb, S. C. (1998). Broadband seismology and noise under the ocean, *Rev. Geophys.* **36,** no. 1, 105–142, doi: 10.1029/97RG02287.

Webb, S. C., and W. C. Crawford (1999). Long-period seismology and deformation under ocean waves, *Bull. Seismol. Soc. Am.* **89,** no. 6, 1535–1542.

Webb, S. C., and W. C. Crawford (2010). Shallow-water broadband OBS seismology, *Bull. Seismol. Soc. Am.* **100,** no. 4, 1770–1778, doi: 10.1785/0120090203.

Yang, Z., A. Sheehan, J. A. Collins, and G. Laske (2012). The character of seafloor ambient noise recorded offshore New Zealand: Results from the MOANA ocean bottom seismic experiment, *Geochem. Geophys. Geosys.* **13,** Q10011, doi: 10.1029/2012GC004201.

Yoshizawa, K., K. Yomogida, and S. Tsuboi (1999). Resolving power of surface wave polarization data for higher-order heterogeneities, *Geophys. J. Int.* **138,** no. 1, 205–220, doi: 10.1046/j.1365-246x.1999.00861.x.

Zha, Y., S. Webb, and W. Menke (2013). Determining the orientations of ocean bottom seismometers using ambient noise correlation, *Geophys. Res. Lett.* **40,** no. 14, 3585–3590, doi: 10.1002/grl.50698.

Zha, Y., S. C. Webb, S. Wei, D. Wiens, D. Blackman, W. Menke, R. Dunn, and J. Conder (2014). Seismological imaging of ridge–arc interaction beneath the Eastern Lau Spreading Center from OBS ambient noise tomography, *Earth Planet. Sci. Lett.* **408,** 194–204, doi: 10.1016/j.epsl.2014.10.019.

Institute of Geophysics and Planetary Physics
Scripps Institution of Oceanography
University of California, San Diego
IGPP-0225
9500 Gilman Drive
La Jolla, California 92093
adoran@ucsd.edu
glaske@ucsd.edu