Geophysical Journal International

Geophys. J. Int. (2021) **224,** 1607–1620 Advance Access publication 2020 November 12 GJI Seismology



Connecting beamforming and kernel-based noise source inversion

Daniel C. Bowden[®],¹ Korbinian Sager[®],² Andreas Fichtner^{®1} and Małgorzata Chmiel^{®3}

¹Department of Earth Sciences, ETH Zürich, 8092 Zürich, Switzerland. E-mail: daniel.bowden@erdw.ethz.ch ²Department of Earth, Environmental and Planetary Sciences, Brown University, 324 Brook St, Providence, RI 02912, USA ³Department of Civil, Environmental and Geomatic Engineering, ETH Zürich, 8093 Zürich, Switzerland

Accepted 2020 November 9. Received 2020 September 17; in original form 2020 June 5

SUMMARY

Beamforming and backprojection methods offer a data-driven approach to image noise sources, but provide no opportunity to account for prior information or iterate through an inversion framework. In contrast, recent methods have been developed to locate ambient noise sources based on cross-correlations between stations and the construction of finite-frequency kernels, allowing for inversions over multiple iterations. These kernel-based approaches show great promise, both in mathematical rigour and in results, but are less physically intuitive and interpretable. Here we show that these apparently two different classes of methods, beamforming and kernel-based inversion, are achieving exactly the same result in certain circumstances. This paper begins with a description of a relatively simple beamforming or backprojection algorithm, and walks through a series of modifications or enhancements. By including a rigorously defined physical model for the distribution of noise sources and therefore synthetic correlation functions, we come to a framework resembling the kernel-based iterative approaches. Given the equivalence of these approaches, both communities can benefit from bridging the gap. For example, inversion frameworks can benefit from the numerous image enhancement tools developed by the beamforming community. Additionally, full-waveform inversion schemes that require a window selection for the comparisons of misfits can more effectively target particular sources through a windowing in a beamform slowness domain, or might directly use beamform heatmaps for the calculation of misfits. We discuss a number of such possibilities for the enhancement of both classes of methods, testing with synthetic models where possible.

Key words: Waveform inversion; Seismic interferometry; Seismic noise; Wave propagation.

1 INTRODUCTION

Recent works have developed a framework to locate ambient noise sources based on cross-correlations between stations and the construction of finite-frequency kernels (e.g. Tromp *et al.* 2010; Hanasoge 2013; Ermert *et al.* 2017; Sager *et al.* 2018; Datta *et al.* 2019; Xu *et al.* 2019), drawing on methods from helioseismology (Gizon & Birch 2002). The approach shows promise, in regards to both the mathematical formalism and in actual long period results.

In this paper, we link this finite-frequency gradient-based inversion approach to the more commonly used beamforming class of methods, such as correlation-based beamforming (e.g. Bucker 1979; Hinich 1979; Fizell 1987; Ruigrok *et al.* 2017), backprojection (Ishii *et al.* 2005; Meng *et al.* 2012), and matched field processing (MFP, Baggeroer *et al.* 1988; Kuperman & Turek 1997). Specifically, certain versions of these algorithms can be shown as mathematically identical to the finite-frequency inversion, for at least the first iteration.

The first half of this paper (Sections 2 and 3) reviews the basic steps of a correlation-based beamforming and MFP algorithm, the latter of which maps sources to a spatial domain instead of assuming plane waves in a slowness or wavenumber domain. From there, a series of modifications to change the order of processing steps and to account for prior information are introduced, ultimately resulting in the kernel-based inversion scheme. Throughout the manuscript, we purposely distinguish and separate those steps which are needed to understand the iteration process from those steps which are used for computational efficiency: adjoint techniques and a direct simulation of correlation wavefields are described for completeness, but only at the end of Section 3, and do not have a direct bearing on the other conclusions of the paper. We emphasize that while this first half of the paper approaches the algorithms in a new and simple way, the actual, final inversion method is unchanged from previous works (Tromp et al. 2010; Datta et al. 2019; Xu et al. 2019, etc.).

Understanding how these two classes of methods relate allows for the enhancement or extension of both. In the second half of the paper, Section 4 demonstrates a few such possibilities. We first

© The Author(s) 2020. Published by Oxford University Press on behalf of The Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

discuss two possibilities relating to the enhancement of beamform techniques: an equivalence between noise-correlations and beamform maps allows for a characterization of the effect of noise preprocessing schemes, and the inclusion of prior information in beamform techniques allows for improved resolution of recovered energy. We then demonstrate two ways in which beamform techniques may benefit a full-waveform inversion scheme, through the use of windowing in slowness space to more intelligently select traveltime windows for misfit calculations, and through the use of beamform heatmaps themselves as misfit functions.

These methods should help improve our ability to image and understand the sources of seismic signals in the Earth. Much of this work is motivated by the desire to understand the distribution of noise sources in the primary and secondary microseism band (e.g. Stehly et al. 2006; Gerstoft & Tanimoto 2007; Juretzek & Hadziioannou 2016), strongest between roughly 5-20 s periods, as the advance of surface wave tomography from ambient noise crosscorrelation has allowed unprecedented crustal imaging in recent years (e.g. Shapiro et al. 2005; Larose et al. 2006). Generally, the technique requires an assumption that the noise wavefield is stationary and this is most likely usually not true, but full-waveform tomographic methods can take a non-homogeneous distribution of sources into account (Sager et al. 2018, 2020). Aside from surface wave tomography, various similar imaging methods are used at higher frequencies to study or exploit anthropogenic noise (e.g. Nakata et al. 2011; Riahi & Gerstoft 2015; Chmiel et al. 2016), small volcano- or tectonic-related tremor (Wassermann 1997; Ghosh et al. 2012; Inbal et al. 2016), the movement of and hydrodynamic systems around glaciers (Walter et al. 2015; Aso et al. 2017; Sergeant et al. 2020), or even to image the complex rupture dynamics of earthquakes (Meng et al. 2012; Gallovic et al. 2019).

Finally, we note that this paper specifically focuses on building a link between these two classes of methods, beamforming/backprojection and kernel-based inversions, both of which exploit correlations between stations. In contrast, direct modelling of individual waveforms from an earthquake source, and the widely adopted migration tools of the seismic exploration community (Song *et al.* 2019), both also offer numerous similarities and advantages [see Li *et al.* (2020) for descriptive overview]. Describing and comparing station-pair correlation methods with direct, singlestation methods may be the scope of future works.

2 TRADITIONAL BEAMFORMING AND MFP

2.1 A description of cross-correlation beamforming

First we review the steps used for a simple beamformer. As will be described shortly, this is based on the cross-correlation of waveforms (Ruigrok *et al.* 2017) rather than the 'delay-and-sum' approach, as the cross-correlation implementation is more directly relatable to later sections. We first assume that observations of a wavefield are collected at an array of stations. A grid search is performed over a range of possible incident slownesses: $s = (s_x, s_y)$, mapping out the possible azimuths from which a plane wave might have originated, and the wave speed that it propagates at. To test whether that particular plane wave is a reasonable description of energy sources, the appropriate station-station time delays are calculated. The time delay between station *i* and station *j* is given as:

$$\Delta t_{ij} = \mathbf{s}^T \big(\mathbf{x}_i - \mathbf{x}_j \big), \tag{1}$$



Figure 1. Calculating time-delays in beamforming: Assuming a possible plane wave (black arrow and wave fronts) incident upon the array (stations at x_1 , x_2 and x_3), time delays are calculated for each station pair.

where x_i and x_j describe the respective coordinates in space. Fig. 1 is a schematic representing three stations and the relative time delays, for a given possible incident plane wave. This assumption of a plane wave allows beamforming algorithms to be extremely computationally efficient, but also represents a significant limitation: the whole approach only works for sources that are well outside the domain of the array and the recorded curvature of the wave front is minimal.

One or the other waveform is shifted, and the correlation coefficient between the two waveforms is calculated:

$$C_{ij}^{obs}(\Delta t_{ij}) = \int u[\mathbf{x}_i, t - \Delta t_{ij}(\mathbf{s})] u(\mathbf{x}_j, t) \,\mathrm{d}t, \qquad (2)$$

where $u(x_i, t)$ refers to the observed time-series, recorded over some time interval, here only on the vertical component of motion. The sum of all station–station pairs' correlation coefficients gives a total power or 'score' for the incident slowness gridpoint being tested:

$$P(s_x, s_y) = \sum_{i,j} C_{ij}^{obs}(\Delta t_{ij}).$$
(3)

We show a synthetic example of this process in Fig. 2. An array is used that is intended to mimic a subset of 11 stations from the Parkfield Array in southern California, with a homogeneous distribution of weak noise sources and a particularly strong noise source to the North East (shown in Fig. 2a). Synthetic seismograms are created using 2-D, analytic, Rayleigh-wave Green's functions, generating roughly four and half hour's worth of noise filled with signals of random phase, as described by Fichtner et al. (2017). These analytic terms account for source-to-station 2-D geometric spreading, a constant Rayleigh wave speed of 3 km s⁻¹, and a constant attenuation of Q = 100. We expect our array is able to resolve wavelengths comparable to the average interstation distance of roughly 6 km. Given the constant model wave speed of 3 km s^{-1} , this translates to a 2-s period. We ultimately use a bandpass filter for our signals between 0.18 and 3.3 s period to cover a range of the various interstation distances present in our sample array. Ultimately, higher or lower frequencies might be resolveable as well; the exact array geometry, widest-point aperture and interstation distances will all play a role in the final array resolution and array response (Rost & Thomas 2002; Gal & Reading 2019). The resulting beamform image is shown in Fig. 2(b), where we see strong arrivals at surface



Figure 2. An example array is shown in panel (a), along with the synthetic noise sources. Each grey pixel represents a randomly distributed background noise source, with a darker, denser, 8-times stronger noise source in the northeast. In panel (b), the resulting beamform image is shown. A small black dot in panel (b) marks the location of the expected plane wave azimuth and slowness.

wave speeds generally, with a stronger hotspot to the northeast as expected.

We note that many other flavours of beamform techniques exist, including more commonly used methods that directly sum timeshifted traces (Bucker 1979), methods that operate in frequency domain and wavenumber space (van Veen & Buckley 1988), some methods which include autocorrelations for energy scaling and those that do not, adaptive methods for higher-resolution recovery such as the Multiple Signal Classification (MUSIC) algorithm (Schmidt 1986) or Multi-rate Adaptive Beamforming (MDVR, Capon 1969), methods which exploit 3-component polarizations (Juretzek & Hadziioannou 2016; Löer et al. 2018) and even neural network or other machine learning applications (Ozanich et al. 2020). Additionally, we note that the approach described above is not the most computationally efficient; it would be smarter to shift all traces to a common reference point rather than loop through every station-station combination. Methods operating in frequency domain may also be more computationally efficient, since a single correlation matrix can be constructed before gridsearching and applying phase shifts. In that case, multiplying in a set of complex phase advances or delays can be compactly expressed as projection onto a set of steering vectors (van Veen & Buckley 1988; Rost & Thomas 2002). While each method will have its own advantages and disadvantages, the time-domain method described in this section is the most directly relatable to what will come in the rest of the paper.

2.2 Extension to backprojection and MFP

Before moving on to finite-frequency kernels, we note that the crosscorrelation beamforming described above may instead be mapped to a spatial domain. That is, not only to describe in which direction the signals are travelling, but where they might have originated. One approach to accomplish this is backprojection (Ishii *et al.* 2005; Walker & Shearer 2009; Meng *et al.* 2012; Fan & Shearer 2015), where generally one would use a beamform heatmap to determine source azimuths (as above), and then shoot rays along a great-circle path, or some other ray-theoretical path derived from a velocity model, to ultimately point back towards some location in space. This has become a popular tool for imaging fault-rupture dynamics, since in this case an endpoint for the backprojected energy is a well-established fault plane. Approaches to locate sources in the



Figure 3. Estimating time-shifts can also be mapped directly onto a spatial domain. Here, a grid search over x-y coordinates facilitates time-shifts that need not assume plane-wave incidence. Even more complex velocity models and ray tracing could be used (Gal & Reading 2018).

subsurface are also regularly used in the seismic exploration literature, and specifically time-reversal migration approaches bear similarity to the others mentioned here (Song *et al.* 2019).

Another variation is MFP (Baggeroer *et al.* 1988; Kuperman & Turek 1997). In this case the time delays, Δt_{ij} , may be defined based on traveltimes from gridpoints in an actual, spatial domain, as in Fig. 3. Compared to eq. (1), we now assume a constant velocity model and estimate the time delay for a given noise source location, x, as:

$$\Delta t_{ij} = \frac{||\boldsymbol{x} - \boldsymbol{x}_i||}{v} - \frac{||\boldsymbol{x} - \boldsymbol{x}_j||}{v}.$$
(4)

Here, symbols ||.|| denote the vector norm for our 2-D example, but 3-D noise sources and station locations could also be used. In this simple case we assume a constant propagation velocity, v, although an arbitrarily complex velocity structure could be accounted for in estimating these Δt_{ij} 's. For example, this could be important when considering the velocity contrast between continental and oceanic crust (Gal & Reading 2018). By not restricting a search to plane waves, the approach is especially useful when sources are relatively close, or even inside the array (Corciulo *et al.* 2012; Chmiel *et al.* 2016; Sergeant *et al.* 2020). Although the MFP algorithm calculates a time-shift, Δt_{ij} , in a different way, the rest of the algorithm remains similar: grid search through the domain, define and apply a given set of time-shifts, and finally measure the resulting correlation scores.

Fig. 4(a) shows an example of this for a two-station array. We show this first for the two-station case since the resulting hyperbolas may be easily interpreted and understood. This uses only a single noise source to the northeast, marked by a red star. If multiple station pairs were present, as in the three-station case of Fig. 4(b), the resulting sum of such hyperbolas better localizes the source.

As with beamforming, we note again that many flavours and variations of such MFP or backprojection techniques exist in literature, sometimes by other names such as the source scanning algorithm by Kao & Shan (2004) or a cross-correlation backprojection by Aso et al. (2017), or by using different methods for estimating the path effect of propagation such as the use of waveform templates by Hawthorne & Ampuero (2017). As with beamforming, many MFP algorithms will operate in frequency domain and apply time-shifts as phase delays or advances only after correlating; the resulting image is theoretically unchanged from the time-domain version. As with beamforming, MFP algorithms might incorporate multiple components and constrain wave types (Gal & Reading 2018; Xu et al. 2019), include various pre-processing steps (Aso et al. 2017), alternatively measure time-shifted envelopes or other metrics (Gal et al. 2019), or use various compressive sensing or Bayesian-formulated regularization to improve the final image (Gemba et al. 2017a, b). The applied time-shifts are more generally referred to as a 'replica vector', which could in principle account for more complexity than a simple time-shift; relative amplitude terms between the stations (Corciulo et al. 2012), or even a full synthetic Green's function might be used to account for additional complexities (discussed further in Section 3.2 and the Appendix).

While we focus on the relatively simple 1-component, constant velocity MFP algorithm described above for the purpose of this paper, we will sometimes consider all variations of the method to be under the same class of algorithms as beamforming: all are primarily data-processing driven (i.e. without an explicit misfit function between observations and synthetics), all are based on the timeshifting observed signals (or otherwise modifying the shape and amplitude), and all finally measure success of that shifting under one metric or another. Casting such algorithms into a single group should not diminish the advantages of a diverse body of research, rather on the contrary it is a motivation to recognize that they all relate; nuanced advantages of a given algorithm might be readily incorporated into others.

3 A DESCRIPTION OF FINITE-FREQUENCY KERNELS

The following is meant to build an intuitive picture of a finitefrequency kernel-based inversion for noise sources, starting from the MFP steps described above. We refer the reader to work by Tromp *et al.* (2010), Hanasoge (2013), Ermert *et al.* (2017) or Sager *et al.* (2018) for a full and mathematically rigorous formulation of the method, to Hanasoge (2013) for a detailed numerical demonstration, and to Datta *et al.* (2019), Ermert *et al.* (2017) and Sager *et al.* (2020) for real-data applications on a local and global scale, respectively. We note that unlike previous derivations, we specifically distinguish between steps that are required for the algorithm to function (described first) and those which are used only for computational efficiency (described later).

3.1 Step one: change the order of operations

The result described above for MFP may be achieved, even if we change the order of processing steps. We might first collect the full noise cross-correlation function (NCF) between each station, $C_{ij}(t)$, drawing on a wide body of literature in ambient noise interferometry (e.g. Shapiro *et al.* 2005; Larose *et al.* 2006):

$$C_{ij}^{\text{obs}}(\Delta t_{ij}) = \int u(\boldsymbol{x}_i, t - \Delta t_{ij}) u(\boldsymbol{x}_j, t) \,\mathrm{d}t,$$
(5)

$$= \int \delta(t - \Delta t_{ij}) \left[\int u(\mathbf{x}_i, \tau - t) u(\mathbf{x}_j, \tau) \, \mathrm{d}\tau \right] \mathrm{d}t, \tag{6}$$

$$= \int \delta(t - \Delta t_{ij}) C_{ij}^{obs}(t) \,\mathrm{d}t, \tag{7}$$

where x_i and x_j denote coordinate vectors of the observations. That is, the measurement is the same whether we apply a time-shift and then measure a zero-lag correlation, or if we compute a crosscorrelation function and then look at a certain lag time. The latter is depicted in Fig. 5, where our noise correlation function is multiplied by an example delta function to pick out a particular time on the causal side. This result might not be surprising to practitioners using MFP or beamform techniques in the frequency domain; the correlation in frequency domain becomes a direct multiplication of the two spectra and a phase shift (the replica vector in this case).

Now we may proceed as before, where a grid search over possible noise sources in our spatial domain is performed. Again, the appropriate time delays are estimated. Technically, eq. (7) could be used for any flavour of beamforming where a time-shift and correlation score are used, but we will remain focused on the spatial domain as used in MFP. We note that this Δt_{ij} is only represented as a delta function; no other information about wave propagation, scattering, amplitude or frequency content is yet given. We also note that the idea of using a 'cross-correlation backprojection' is implemented by Aso *et al.* (2017), which also exploits this similarity in processing steps.

The sum over all station-station pairs again gives the power for a noise source at that location in space:

$$P(x, y) = \sum_{i,j} \int \delta(t - \Delta t_{ij}(x, y)) C_{ij}^{obs}(t) dt.$$
(8)

Here, we change the notation from a bold vector **x** for spatial location to an explicit 2-D coordinate *x* and *y* to emphasize that we are gridsearching through the spatial domain. The application of eq. (8) results in a source localization identical to the MFP result shown previously in Fig. 4. Rearranging terms in this way is relatively trivial. It is a key step in linking MFP and an iterative inversion framework, however, since both approaches operate on $C_{ij}^{obs}(t)$. From here we describe a set of enhancements or modifications to eq. (8), leading to the iterative, kernel-based inversion algorithm.

3.2 Modification: using a more specific window function

The delta function used in eq. (8) can be replaced with something more generic. We can replace the delta function with a term W_{ij} :

$$P(x, y) = \sum_{i,j} \int W_{ij}(t, x, y) C_{ij}^{obs}(t) dt.$$
 (9)

Here we note that our noise-source inversion framework will first start to diverge from MFP. As described in Section 2.2, such a term as W_{ij} takes the place of a replica vector in MFP, and might represent more than just a time-shift for the direct arrival. It might include



Figure 4. Synthetic MFP example with two stations (panel a) and three stations (panel b). This uses only one synthetic source in the northeast (red star), and a fewer number of stations, so the similarity to sensitivity kernels can later be more clearly seen.



Figure 5. One noise cross correlation and the appropriate delta-function time pick. This shows the two separate pieces in eq. (7), where any precomputed noise correlation function can be used in beamforming or MFP.

an amplitude term (Corciulo *et al.* 2012), discussed further in the Appendix, or it might even be a complete Green's function capable of modelling various wave types or arrivals in a complex velocity model. The goal in any case is to describe, and thus correct for any known physics that might affect the waveforms before the final correlation score is measured.

In contrast, in the noise-source inversion framework that we are moving towards, such synthetic Green's function modelling is introduced in a different way, as an explicit misfit function shown in the next section. In that case, the main purpose of W_{ij} will be to define a finite window over which correlation scores should be measured. This windowing function is a natural extension of the delta function previously used; given that our data is usually bandpassed to a finite range of frequencies, a window of at least the minimum observable wavelength might be used. Alternatively, in the case of existing works using kernel-based inversion schemes as by Ermert *et al.* (2017) or Sager *et al.* (2020), a windowing term was defined to specifically encompass expected surface wave arrivals, since that was the target of the modelling efforts.

3.3 Modification: include prior information

MFP offers no way to iterate, or to compare a model result back to the data that generated it. The same may be said of plane-wave beamforming, which is perhaps why sophisticated imaging enhancement techniques are desirable—only one image may be produced from any set of data, and so any efforts to improve that image's resolution should be taken.

Here we may change this limitation. With some model of where the sources might be, a synthetic wavefield may be calculated and cross-correlations collected for these synthetic waveforms. A misfit can be defined:

$$\mathcal{X}(x, y) = \sum_{i,j} \int \left[W_{ij}(t, x, y) C_{ij}^{\text{obs}}(t) - W_{ij}(t, x, y) C_{ij}^{\text{syn}}(t) \right]^2 \mathrm{d}t.$$
(10)

Here we use a least-squares misfit for simplicity, but other misfit functions may be explored. The misfit used by Ermert *et al.* (2017) is a measurement of energy asymmetry between the causal and acausal branches in the noise-correlation function, rather than a wiggle-for-wiggle time-domain correlation score. We note also that the definition of a 'physical model' for calculating synthetics may include a range of possibilities. The obvious choice is to map noise sources on some physical domain as is done in MFP. Alternatively, the synthetic model may state something about the content of wave types (i.e. a given ratio between surface wave energy and body-wave energy), or test the stationarity of the ambient noise field. The only requirement is that one has a mechanism for generating synthetics according to that hypothesis.

3.4 On gradients and iterating

Most importantly, eq. (10) offers a path forward to iterate. However, first our interpretation of the reason for our grid search in space needs to change. In the traditional MFP framework, the grid search gave us a method by which to define the time-shifts needed for correlation. Now, instead, we must consider our grid search as actually perturbing the strength of a source; we must actually place (or strengthen) a source at that location in order to calculate synthetics for comparison in eq. (10). Whether the misfit increases or decreases indicates whether we should in fact increase or decrease the strength of a source at this location. This spatially maps out our gradient, which in seismology is often also referred to as a sensitivity kernel in this first iteration. Strictly speaking, a gradient and sensitivity kernel need not be defined on the same domain. In any case, we have a measure of how to update our model, and after applying that change we may repeat the process as many times as desired.

An alternate interpretation of existing MFP algorithms is also possible in the framework of eq. (10), first pointed out by Xu *et al.* (2019) in their appendix C. If one considered a starting model with absolutely no noise sources in the domain, the synthetic noise correlations, C_{ij}^{syn} , will all be zero, and we are reduced to eq. (9). In this hypothetical one-step-only iteration scheme, we update our model

according to the computed gradient, and our final model or interpretation ends there. Fig. 6 shows this final model. We note that it differs from Fig. 4 in that an amplitude decay is observed because injecting synthetic source perturbations at each gridpoint results in synthetic seismograms subject to attenuation and geometric spreading, but in all other regards the results here are identical to that of the MFP results in Fig. 4. We finally note, again, that it would have been possible to incorporate amplitude information into an MFP algorithm, further making Figs 4 and 6 more similar, and discuss this further in the Appendix.

Perturbing the source model and actually computing synthetic seismograms offers other advantages. We already mentioned that attenuation and geometric spreading will be taken into account here, but even more realistic wave propagation effects may also be observed if the velocity model used is complex enough. Effects like frequency dispersion of surface waves, multipathing, and scattering or reflections off of heterogeneities will all give complexities in the synthetics that ideally will match the observed seismograms. That is, in beamforming and MFP, such an effect as a reflection or scatterer might often be misinterpreted as an additional unique source (Ma et al. 2013), whereas here the waveforms would only optimally match if both the direct arrival and some secondary arrival were present in the synthetic waveform. Of course this assumes one has already developed a good velocity model for a given region. The application of an iterative scheme similar to eq. (10) to map noise sources, alternating with full-waveform inversions to improve the velocity structure, was precisely the goal described by Sager et al. (2018).

In some sense, our comparison of the methods may be concluded here as we have described a fully functioning inversion algorithm. However, the approach described above would be quite computationally expensive. At least for ambient noise sources, the approach would require generating synthetic noise sources at every location in a spatial domain, separately, and for a long enough time window to gain a stable estimate of the noise correlation functions. To find a more efficient way, the approach of Tromp *et al.* (2010) directly calculates the correlation wavefield and also uses adjoint techniques. These are both described below, but we emphasize neither is strictly needed to understand how an inversion might be performed.

3.5 Modification: calculate a synthetic correlation wavefield for more efficiency

Generating synthetics for ambient noise cross-correlations can be prohibitively expensive. An impulsive earthquake source can be modelled with only the time steps needed for waves to propagate across the domain, but generally computing NCFs requires averaging time-series over multiple hours or days, and thus multiple realizations of the random noise field. This section describes a more efficient approach.

We start by writing the time-averaged correlation that we would recover through traditional means, $\langle C_{12} \rangle$. Following the approach and workflow developed for active sources (depicted by the cartoon in Fig. 7a), a random noise source $N(\mathbf{x})$ is injected at \mathbf{x}_s and the resulting noise field is recorded at the stations. According to the representation theorem, the seismograms $u(\mathbf{x})$ are in the frequency domain given by:

$$u(\mathbf{x}_1) = G(\mathbf{x}_1, \mathbf{x}_s) N(\mathbf{x}_s), \tag{11}$$

$$u(\mathbf{x}_2) = G(\mathbf{x}_2, \mathbf{x}_s) N(\mathbf{x}_s), \tag{12}$$

where $G(\mathbf{x}, \mathbf{x}_s)$ denotes the Green's function at \mathbf{x} due to an impulse in time acting at \mathbf{x}_s . For the sake of simplicity, we only present equations for scalar wavefields and for two stations. One realization C_{12} of the cross-correlation can then be computed according to:

$$C_{12} = u^*(\mathbf{x}_1)u(\mathbf{x}_2) = G^*(\mathbf{x}_1, \mathbf{x}_s)G(\mathbf{x}_2, \mathbf{x}_s)N^*(\mathbf{x}_s)N(\mathbf{x}_s),$$
(13)

where * denotes complex conjugation. In order to obtain the ensemble average C_{12} , we repeat the whole procedure many times with different realizations of the random noise sources and obtain:

$$\langle C_{12} \rangle = G^*(\mathbf{x}_1, \mathbf{x}_s) G(\mathbf{x}_2, \mathbf{x}_s) \langle N^*(\mathbf{x}_s) N(\mathbf{x}_s) \rangle, \qquad (14)$$

where $\langle \cdot \rangle$ indicates ensemble averaging. In eq. (14) we assume that changes of the Green's functions in the course of the realizations are negligible.

Since the Green's functions do not change, using analytical solutions, or the use of a database of precomputed Green's functions can save significant computation time. In addition, source–receiver reciprocity can be exploited if the number of noise sources is larger than the number of receivers. Although the computational burden is thus reduced, we still have to compute multiple realizations. This is typically only done for demonstration purposes (van Driel *et al.* 2015) or to advance our understanding of the underlying physics (Cupillard & Capdeville 2010; Fichtner *et al.* 2017).

Recent studies (i.e. Tromp *et al.* 2010) offer a more efficient alternative, drawing originally from the field of helioseismology (Gizon & Birch 2002). We may skip the computation of raw, single-station time-series and instead directly target the ensemble averaged correlation function:

$$\langle C_{12} \rangle = G^*(\mathbf{x}_1, \mathbf{x}_s) G(\mathbf{x}_2, \mathbf{x}_s) S(\mathbf{x}_s), \tag{15}$$

where the term $S(\mathbf{x}_s)$ is equivalent to the ensemble average $\langle N^*(\mathbf{x}_s)N(\mathbf{x}_s)\rangle$, that we ultimately want to describe spatially. In this case we do not need multiple realizations or long time periods to average; we directly describe a source power spectrum, $S(\mathbf{x}_s)$, everywhere in space.

The steps to to directly model $S(\mathbf{x}_s)$ according to eq. (15) are described below:

(i) Choose a reference station, for example at \mathbf{x}_1 in Fig. 7(b), and compute the Green's function, $G(\mathbf{x}_s, \mathbf{x}_1)$. That is, inject a delta impulse in time at \mathbf{x}_1 , as in Fig. 7(b).

(ii) Assuming source–receiver reciprocity, time reverse the Green's function and multiply it with the assumed power-spectral density distribution at a given source location: $S(\mathbf{x}_s)$. We thus have $G^*(\mathbf{x}_1, \mathbf{x}_s)S(\mathbf{x}_s)$.

(iii) Inject the result of the previous step as a source into a numerical solver and simulate the wavefield. Sample the wavefield where other receivers are located, in our example case at \mathbf{x}_2 in Fig. 7(c). Following the representation theorem, the recording is then given by $G(\mathbf{x}_2, \mathbf{x}_s)[G^*(\mathbf{x}_1, \mathbf{x}_s)S(\mathbf{x}_s)]$, which is equivalent to the ensemble averaged correlation function in eq. (14).

With this recipe, we obtain a set of ensemble averaged correlation functions that have a common reference station, only requiring 2 wavefield simulations, again only requiring the computation time needed for one set of signals to propagate across the domain. We can directly model the power-spectral density, without needing multiple realizations or a very long time-series. The result is the ideal NCF that many traditional realizations of the noise field would have converged to.

The cartoons in Fig. 7 only illustrate an example with one noise source. Eq. (14) can be generalized by adding an integral over



Figure 6. Synthetic sensitivity kernels with two stations (panel a) and three stations (panel b). As before, there is only one noise source in the northeast, indicated with a red star. These kernels have the same shape and structure as the MFP results before, only the amplitude decays because of how synthetic waveforms are analytically modelled.



Figure 7. A brute-force scheme is depicted in panel (a), where a given noise source is placed, waves recorded at two stations, and these are then cross-correlated over many realizations of the wavefield (denoted by the angle brackets). This is representative of eq. (14). Panels (b) and (c) together depict the correlation wavefield approach, described in text and eq. (15).

noise sources. The implicit assumption then is that neighbouring noise sources are uncorrelated, a common assumption also invoked for the principal of Green's function retrieval (Wapenaar 2004; Wapenaar & Fokkema 2006). While an approach using precomputed noise correlations might simply sum multiple sources in Fig. 7(a), all non-zero sources in panel (c) would act simultaneously.

3.6 Modification: adjoint techniques for more efficiency

We now have a method for efficiently calculating NCFs, for any arbitrary model of sources. Now the question for our gradientbased inversion is: how does a perturbation to a given individual noise source affect the misfit between observed NCFs and synthetic NCFs. In the most straightforward approach, we would only perturb one source at a time to measure its effect, and again find ourselves in a situation requiring simulations for every gridpoint.

The solution here is to exploit adjoint techniques in order to find efficient expression for the computation of gradients. In the following, we circumvent mathematical derivations and show how the same expressions can be obtained based on heuristic arguments. Previously, we showed that a gridpoint in space can be connected to a particular lag-time location on the NCF. Now we consider the other direction: that a particular time in the NCF will correspond to a range of possible locations in space. For zero-lag, for example, we may initiate a simulation from both stations, simultaneously, and record where the two wavefields interact. This is shown in the top row of Fig. 8, where the grey shaded region records every point that can possibly contribute to the zero-lag time of the NCF between these two stations. This zero-lag source kernel can be written as:

$$K(\mathbf{x}) = \int G(\mathbf{x}, \mathbf{x}_1, t) G(\mathbf{x}, \mathbf{x}_2, t) dt.$$
 (16)

That is, although we simulated two Green's functions, $G(x, x_1, t)$ and $G(x, x_2, t)$, we assume that the sources could have instead been initiated elsewhere in space and recorded at x_1 and x_2 , that is we invoke reciprocity for both receivers.

If, instead, we are interested in the causal surface wave traveltime, $\Delta t_{\rm SW}$, we can start the injection of the wavefield at the receiver (in the sense of the modelling recipe above) at earlier times, at $-\Delta t_{\rm SW}$, while the wavefield at the reference station still starts at t = 0. The region of interaction then maps out the stationary zone for surface waves behind the second station, as in the bottom row of Fig. 8. In this case the kernel can be written as:

$$K(\mathbf{x}) = \int G(\mathbf{x}, \mathbf{x}_1, t) G(\mathbf{x}, \mathbf{x}_2, t + \Delta t_{\rm SW}) \,\mathrm{d}t.$$
(17)

This is still missing two points, however. We previously stated that we do not want merely a delta-function time-shift in the data, but that we: (1) want to consider a more general window of the NCF (Section 3.2) and (2) now are measuring a misfit between an observed NCF and a simulated NCF, and thus want to know what the amplitude of that misfit implies about the perturbation (Section 3.3). We can achieve both points by replacing the delta function with a source time function, f(t), to be injected into the wavefield simulations above.

This f(t) solves the first problem above by being defined over some finite length around the desired propagation time. To achieve the second goal and measure what a given misfit implies about a possible



Figure 8. Schematics for adjoint wavefields: Panels (a)–(e) depict successive time steps for the case of correlation energy at zero-lag time, and panels (f)–(j) depict time steps for the case of energy in the surface wave time window. In both examples, the two circles (yellow and blue) represent the wave front coming from either station, where they are currently interacting is highlighted in red, and the region mapped out by that interaction is indicated in grey.

perturbation to the model, we can define f(t) as a derivative of the misfit function with respect to the synthetic correlation function. It then directly carries information about how the increase or decrease of a given $C_{ij}^{syn}(t)$ will increase or decrease that misfit. Now, we can give this generic function, f(t), its proper name: the adjoint source. In the case of a least-squares misfit from eq. (10), for our two station example, we are left with:

$$f(t) = 2 \cdot [W_{12}(t)C_{12}^{obs}(t) - W_{12}(t)C_{12}^{syn}(t)],$$
(18)

$$K(\mathbf{x}) = \int f(-t) \star G(\mathbf{x}, \mathbf{x}_2, t) G(\mathbf{x}, \mathbf{x}_1, t) dt, \qquad (19)$$

where \star denotes a convolution, and we note that eq. (19) includes a time-reversal of f(t). Recordings at the beginning of the simulation, for instance, have not travelled far and they are thus only influenced by noise sources close to the station. This also justifies the convention used in Fig. 8(f), where the causal branch of surface wave signals on the NCF are recovered by simulating the right Green's function starting at negative times.

We now have an expression to compute a source sensitivity kernel that indicates how and which noise sources in space change our misfit function. The description above does not constitute a rigorous proof, and again we refer the reader to work by Tromp *et al.* (2010) or others. We also emphasize that adjoint techniques are, generally speaking, a broader mathematical tool generalizing the transpose of a matrix. The interpretation of the adjoint wavefield as a physical wavefield is not always possible nor strictly necessary (Thrastarson *et al.* 2020; van Driel *et al.* 2020); the motivation and physics described in this section are given rather as an after-the-fact interpretation.

The goal of this section was not to reinvent adjoint methods or a new inversion algorithm. The beamforming algorithm, MFP, and kernel-based noise source inversion have all been formulated previously in other works. Rather, pointing out the similarity between a purely data-driven algorithm like MFP or beamforming (eq. 9) and the gradient-based inversion approach (eq. 10) opens the doors to multiple ways either class of method may be improved.

4 COMBINING APPROACHES, INSTEAD OF JUST COMPARING

The similarities between beamforming/MFP methods and kernelbased noise inversions suggest numerous possibilities for the improvement of either method. In this section, we first emphasize how beamform images can be immediately applied to any noise interferometry study, and illustrate how beamforms can help one understand various noise pre-processing strategies. Secondly, given that the kernel-based inversion framework above allows for the incorporation of a starting model or prior knowledge, we show how beamform heatmaps can be improved by accounting for this knowledge. Thirdly, for the improvement of existing full-waveform noise source inversion schemes, we show how beamform maps allow for a more targeted and more effective selection of misfit windows. Finally, we suggest that a sum over slownesses in beamform heatmaps themselves may be used as a misfit for gradient-based inversion schemes, allowing for the inclusion of more sophisticated techniques to increase resolution (e.g. MUSIC).

Sophisticated MFP image enhancing approaches (e.g. Gemba *et al.* 2017b) might also be incorporated into the kernel-based framework. While that might prove the most fruitful line of eventual research, the suggestions here relate rather to the improvement and incorporation of beamforming methods due to their low computational cost. As mentioned, beamforming generally requires assuming signals are plane waves and thus distant from the array, meaning their applicability is more limited than MFP or kernel-based noise inversions.

4.1 Producing beamforms from noise-correlations

The fact that pre-computed noise correlation functions can be used to construct beamform maps, as shown in Section 3.1, implies that both views of the data can always be taken where one is already considered. In any ambient-noise interferometry study where station pairs are cross-correlated, it would be computationally trivial to construct a beamform image as well, as is also emphasized by the work of Seydoux *et al.* (2017). It is almost always an assumption in noise-interferometry works that the wavefield is either perfectly stationary, or at least scattered enough that reasonable traveltime estimates can be extracted (Lobkis & Weaver 2001; Tsai 2010). Some authors have explicitly shown beamform estimates to demonstrate this either is the case (e.g. Stehly *et al.* 2006; Seats *et al.* 2012; Zhang & Yang 2013; Spica *et al.* 2018) or is not (e.g. Lehujeur *et al.* 2017), but this is not the norm.

Further, given that these beamform maps can be constructed from noise correlation functions, one may include the effect of any pre- or post-processing, as is also shown by Seydoux *et al.* (2017) and Girard & Shragge (2020) for beamforms, and by Aso *et al.* (2017) for backprojection. Generally it is believed that signal processing of either the individual noise pairs or the final NCF will reduce the effect of non-stationary, impulsive sources (Bensen *et al.* 2007). Fig. 9(a) shows our same raw beamform as before, and a



Figure 9. Using processed waveforms: one-bit normalization, spectral whitening and averaging of the causal and acausal branches.

version of the beamform constructed from NCF's that include various processing. Panel (b) uses data where one-bit normalization was applied, and we find for this synthetic data that it actually does very little to improve the stationarity of the noise wavefield. Panel (c) uses spectral whitening, and while we do see that it reduces the effect of the strong noise source in the northeast, it is also more sharply pixelated, implying we have potentially decreased the stationarity at certain other slownesses. Panel (d) uses an averaging of the causal and acausal branches; while the relative strength of the noise source in the northeast is somewhat reduced, we also find a mirrored version of the strong source to the southwest, which might not be desirable. We also refer readers to work by Fichtner et al. (2017) for an example of how these processing methods can affect source or structure spatially. Finally, we emphasize that the data used here are entirely synthetic and the application of various pre-processing strategies to real data may look quite different; the point is rather that such maps can be used to understand the effects of these various processing schemes on any data set.

We thus advocate that beamform maps should be regularly included in noise interferometry studies, and further that any beamform map should include identical pre-processing and station-pair selection as the interferometry. If certain station pairs are removed due to low signal-to-noise ratios (SNR), this can also be reflected in the accompanying beamform image. Such pre-processing or SNR criteria are often well justified; our point here is that with a database of processed noise correlations, it requires virtually no additional computational effort to produce a corresponding beamform image. Similarly, if it was believed that noise sources may be present very near to the array or within the array bounds, an MFP algorithm could also use these precomputed and processed noise correlations.

4.2 Priors on beamforming

We previously discussed how the standard MFP algorithm may be considered analogous to the more complete gradient-based inversion scheme, mainly by being explicit about the fact that an MFP starting model contains no noise sources. The same may be said of beamforming in slowness or wavenumber domain, assuming a starting model of zero energy everywhere. Eq. (20) below is analogous to eq. (10) earlier, where we calculate waveform misfits between the data correlations and our model correlations, but now we consider possible incident slownesses:

$$\mathcal{X}(s_x, s_y) = \sum_{i,j} \int \left[W_{ij}(t, s_x, s_y) C_{ij}^{\text{obs}}(t) - W_{ij}(t, s_x, s_y) C_{ij}^{\text{syn}}(t) \right]^2 dt.$$
(20)

We see that if a synthetic starting model consisted of zero noise sources, and thus zero correlation in each $C_{ij}^{syn}(t)$, this misfit reduces to the same as a standard correlation beamformer (e.g. eq. 9). However, assuming a prior model consisting of absolutely no noise sources is actually a rather strong statement, given that one is generally using a beamforming method specifically to search for noise sources.

Rather, what most researchers in the interferometry community often assume is the presence of isotropic and stationary noise sources distributed across the Earth's surface (Aki 1957; Boschi & Weemstra 2015). This is a contradiction—that researchers would assume isotropic noise sources for one method and implicitly assume zero noise sources for another.

To account for isotropic noise sources on the Earth's surface as a prior model, we can specify terms in eq. (20). C_{ij}^{obs} is our observed data correlations (in our case synthetically generated for the example), and the beamform resulting from this is shown in Fig. 10(a). Here, the noise sources consist of a randomly distributed sample of weak noise sources, with a stronger noise source to the northeast, as in our original Fig. 2(a). The term C_{ii}^{syn} will hold our prior knowledge or assumption of a homogenous background without the knowledge of the stronger source. That is, we compute synthetics for a noise wavefield specifically without any knowledge of a stronger source imprinted; the resulting beamform is shown in panel (b), where we see only an even ring of sources arriving from all directions. The misfit between the data and our prior model, calculated as in eq. (20), is shown in panel (c). Even without further iterations, being explicit about assumed prior knowledge is already greatly beneficial.

Normally, under an iterative scheme, one would take this misfit and the knowledge gained, update our model, and repeat the process. In this case we do not yet have a method by which to take our beamform map and return to the noise sources on some spatial domain; the application of eq. (20) works only for this first iteration. Works by Yao & van der Hilst (2009) and more recently Lehujeur *et al.* (2017), however, do propose a framework in which to forward model noise correlations directly from a beamform-style map, and the incorporation of such a method to eq. (20) could be a topic of future work.

4.3 Intelligent windowing

With an explicit connection between beamforming and gradientbased inversions, we can use beamforming as a tool to steer and guide more sophisticated misfit definitions. Current inversion schemes (e.g. Ermert *et al.* 2017; Sager *et al.* 2020) often restrict



Figure 10. Panel (a) shows the beamform image directly from observations as in Fig. 2(b). Using a randomly generated, isotropic distribution of ambient noise sources will result in a ring of slownesses as in panel (b). If we used those NCFs as prior knowledge and subtracted from the observed correlations we are left with an improved recovery, shown in panel (c).

the misfit computation to a time-window of the NCF near the expected arrival time of surface waves, as in the red-shaded regions of Fig. 11(b). This gives a reasonable result for station–station pairs oriented in such a way that they were pointing towards the noise-source of interest, as in 'Station Pair 1' in Figs 11(a) and (b). This windowing ignores other spurious signals in the correlation function, however. If a noise source is oriented such that the signal might arrive at both stations simultaneously (i.e. near zero-lag time, as in 'Station Pair 2', panel c), this information is simply ignored and lost. The final beamform image in Fig. 11(d) only uses information from these restricted time windows, and as a result the localization of the noise source is poorer.

This could be used to our advantage, however, if a particular direction or beamform 'hotspot' is known to be of particular interest. There may be physical situations, for example, where noise sources are expected in a particular ocean basin, and thus only in a particular direction relative to the array. Defining an azimuth range and slowness will restrict the possible windows over which a correlation score is computed, thus improving the misfit estimates in iteration by reducing effects from aliasing, scatterers, and potentially other effects. Figs 11(e)–(g) illustrates this, where a box is drawn around a target slowness, and this box is used to more intelligently pick time windows for later use as a misfit in a finite-frequency inversion scheme. As expected, the strongest signals in each NCF fall within the blue shaded regions.

This idea relates to work by Retailleau *et al.* (2017), who observed consistent spurious arrivals in the noise-correlation functions. They found that these can be used to image sources 'off-axis' of the two arrays of stations used. By gridsearching over a spatial domain covering the Atlantic ocean, they estimated the expected time windows in the correlation functions and compute a more effective slant-stack. What we propose here is similar: by first performing a beamform of raw data, a dominant noise source will be more easily noted and then researchers can consciously choose to include or exclude that source in misfit definitions. This suggested approach of using beamforms is somewhat restrictive however, assuming all sources are outside the array and signals are simple plane waves, whereas the work of Retailleau *et al.* (2017) remains more flexible and can image sources between any given station pair.

We also note that work by Seydoux *et al.* (2017) suggests a similar point, though for different purposes. They have the goal of improving ambient-noise correlation functions for interferometry, and point out that first computing a beamform is beneficial. They show that conditioning of the beamform image, achieved by normalizing eigenvalues of the coherence matrix to display a more isotropic distribution of noise sources, also results in individual noise-correlations functions that are cleaner and easier to interpret traveltimes from. In this case, rather than using windows of the beamform image only to guide where one should look in a correlation function, the correlation functions are themselves modified in some intelligent (though non-linear) way.

4.4 Beamforms as misfits

Another future approach may be to use the beamform image itself as a misfit. All of the iteration schemes described thus far compared NCF waveforms on a station-pair to station-pair basis, using something like a least-squares misfit (eq. 10) or the asymmetry score of Ermert *et al.* (2017). The difference between a synthetically produced beamform map and an observed beamform map is another possible misfit:

$$\mathcal{X} = \sum_{s_x, s_y} \left[\sum_{i,j} \int W_{ij}(t, s_x, s_y) C_{ij}^{\text{obs}}(t) \, \mathrm{d}t - \sum_{i,j} \int W_{ij}(t, s_x, s_y) C_{ij}^{\text{syn}}(t) \, \mathrm{d}t \right]^2.$$
(21)

A potential advantage of this approach comes how aliased features can be interpreted. Aliasing is a common problem both in full-waveform correlation misfits, where a 2π cycle shift can give a high correlation score, and in beamform images where multiple hotspots may appear for the same reason. At least with the creation of beamform images, however, it may be possible to either clean and thus better constrain image features (via MUSIC, for example), to incorporate constraints for wave type using 3-component observations, or to more intelligently select a region of the beamform image to be concerned with (as in Fig. 11). Other image comparison tools also exist that can score similarity based on common features, even if the individual pixels of one image do not perfectly line up with those of another (Wang et al. 2004; Araya-Polo et al. 2018). We note that in this case the interpretation of beamform heatmaps as plane waves is not important; even if a source is relatively close to the array and the plane-wave assumption is violated, we are anyways



Figure 11. Panel a depicts our array of stations, showing the direction of the dominant noise energy (black arrow) and highlighting two station-station pairs. Panels (b) and (c) show how a typical inversion scheme might be limited to surface wave windows, and the resulting degraded beamform image with this restriction is shown in panel (d). No beamformer to our knowledge operates this way, but it demonstrates how data would be lost if windows were too restrictive. In contrast, if we instead use full waveforms to beamform and select a hotspot of interest in the white dashed box in panel (e), we can more effectively window energy in panels (f) and (g).

only using the beamform heatmap as a transformation of the raw time-series.

Additionally, the weighting of information from a beamform image may be different: for current inversion schemes where a misfit is based on waveform differences, all parts of the used waveforms are considered equal. In a typical beamform image, however, pixels near the center (near-zero slowness space) represent more significant changes in incident velocity as compared to pixels near the edges. This may or may not be desirable, depending on the problem of interest.

5 CONCLUSIONS

We show in Sections 2 and 3 that a correlation-based beamforming, backprojection or MFP algorithm can be considered very similar to a gradient-based inversion scheme. The primary difference between these two classes of methods comes in the generation of an explicit prior starting model (or lack thereof). For many authors and applications, it may be that the purely data-driven versions of beamforming and MFP are computationally superior and easier to implement, but at least understanding the connection should be useful towards the development of better algorithms. Again, the beamforming and backprojection communities have spent considerable effort exploring image enhancement techniques (e.g. MUSIC, MDVR), and so it follows that such methods might be incorporated into kernel-based inversion schemes in the near future.

The figures and equations presented here exploited only the vertical component of motion, but the use of horizontal motions has been explored by each of the subdisciplines discussed here: for example in beamforming by Juretzek & Hadziioannou (2016), in Matched-Field Processing by Gal & Reading (2018) and even in finite-frequency source inversions by Xu *et al.* (2019). If anything, the fact that each of these three research communities have separately developed multicomponent frameworks (with relatively little citations across disciplines) further emphasizes the importance of understanding the connection between them.

Section 4 specifically outlines several ways in which this connection may be immediately useful. Beamformer algorithms can be adapted to directly use noise-correlation functions for both computational efficiency and for the purpose of incorporating or testing various pre-processing methods. Beamforming algorithms might also benefit from the explicit definition of a prior starting model. Going in the other direction, a standard beamform that uses all possible data (e.g. looks in all possible directions) might be used to more selectively window noise-correlation functions. Finally, we point out that beamform heatmaps themselves might be used directly as misfits, though further testing and demonstrations with real data will be the subject of future work.

ACKNOWLEDGEMENTS

The authors would like to thank Victor Tsai for helpful discussions and suggestions, and both Nori Nakata and Dylan Mikesell for insightful and helpful reviews. DCB acknowledges support from the Swiss National Science Foundation (SNSF) under the Spark Program grant CRSK-2_190837. KS acknowledges support from the SNSF under grant P2EZP2_184379. Example Python Jupyter notebooks to accompany the paper are online at https://github.com/dcbowden/BeamKer, and on our group software page at https://cos.ethz.ch/software/research.html. Synthetic data used throughout the paper are generated using scripts from the Generealized Interferometry package from Fichtner et al. (2017), also available online at https://cos.ethz.ch/software/research.html.

REFERENCES

- Aki, K., 1957. Space and time spectra of stationary stochastic waves, with special reference to microtremors., *Bull. Earthq. Res. Inst.*, **35**, 415–457. Araya-Polo, M., Jennings, J., Adler, A. & Dahlke, T., 2018. Deep-learning
- tomography, *Leading Edge*, **37**(1), 58–66.
- Aso, N., Tsai, V.C., Schoof, C., Flowers, G.E., Whiteford, A. & Rada, C., 2017. Seismologically observed spatiotemporal drainage activity at moulins, *J. geophys. Res.*, **122**, 9095–9108.
- Baggeroer, A.B., Kuperman, W.A. & Schmidt, H., 1988. Matched field processing: source localization in correlated noise as an optimum parameter estimation problem, *J. acoust. Soc. Am.*, 83, 571–587.
- Bensen, G.D., Ritzwoller, M.H., Barmin, M.P., Levshin, A.L., Lin, F., Moschetti, M.P., Shapiro, N.M. & Yang, Y., 2007. Processing seismic ambient noise data to obtain reliable broad-band surface wave dispersion measurements, *Geophys. J. Int.*, 169(3), 1239–1260.
- Boschi, L. & Weemstra, C., 2015. Reviews of geophysics stationary-phase integrals in the cross correlation, *Rev. Geophys.*, 53, 411–451.
- Bucker, H.P., 1979. Use of calculated sound fields and matched-field detection to locate sound sources in shallow water, *J. acoust. Soc. Am.*, **59**, 368–373.
- Capon, J., 1969. High resolution frequency-wavenumber spectrum analysis, *Proc. IEEE*, 57, 1408–1418.
- Chmiel, M., Roux, P. & Bardainne, T., 2016. Extraction of phase and group velocities from ambient surface noise in a patch-array configuration, *Geophysics*, 81(6), KS231–KS240.
- Corciulo, M., Roux, P., Campillo, M., Dubucq, D. & Kuperman, W.A., 2012. Multiscale matched-field processing for noise-source localization in exploration geophysics, *Geophysics*, 77(5), doi.org/10.1190/geo2011-0438.1.
- Cupillard, P. & Capdeville, Y., 2010. On the amplitude of surface waves obtained by noise correlation and the capability to recover the attenuation: a numerical approach, *Geophys. J. Int.*, **181**, 1687–1700.
- Datta, A., Hanasoge, S. & Goudswaard, J., 2019. Finite-frequency inversion of cross-correlation amplitudes for ambient noise source directivity estimation, J. geophys. Res., 124(7), 6653–6665.
- Ermert, L., Sager, K., Afanasiev, M., Boehm, C. & Fichtner, A., 2017. Ambient seismic source inversion in a heterogeneous Earth: theory and application to the Earth's Hum, *J. geophys. Res.*, **122**(11), 9184–9207.
- Fan, W. & Shearer, P.M., 2015. Detailed rupture imaging of the 25 April 2015 Nepal earthquake using teleseismic P waves, *Geophys. Res. Lett.*, 42(14), 5744–5752.
- Fichtner, A., Stehly, L., Ermert, L. & Boehm, C., 2017. Generalized interferometry – I: theory for interstation correlations, *Geophys. J. Int.*, 208(2), 603.

- Fizell, R.G., 1987. Application of high-resolution processing to range and depth estimation using ambiguity function methods, *J. acoust. Soc. Am.*, 82, 606–613.
- Gal, M. & Reading, A.M., 2018. Matched field processing of threecomponent seismic array data applied to Rayleigh and Love microseisms, *J. geophys. Res.*, **123**, 6871–6889.
- Gal, M. & Reading, A.M., 2019. Beamforming and polarisation analysis, in Seismic Ambient Noise, Chapter 2, pp. 32–72, eds Nakata, N., ,Fichtner, A. & Gualtieri, L., Cambridge Univ. Press.
- Gal, M., Reading, A.M. & Ellingsen, S.P., 2019. Short timescale analysis of microseisms and application to array calibration, *J. geophys. Res.*, 124, 2684–2701.
- Gallovic, F., Valentova, L., Ampuero, J.-P. & Gabriel, A.-A., 2019. Bayesian dynamic finite – fault inversion: 1. Method and synthetic test, *J. geophys. Res.*, **124**, 6949–6969.
- Gemba, K.L., Hodgkiss, W.S. & Gerstoft, P., 2017a. Adaptive and compressive matched field processing, J. acoust. Soc. Am., 141(1), 92–103.
- Gemba, K.L., Nannuru, S., Gerstoff, P. & Hodgkiss, W.S., 2017b. Multifrequency sparse Bayesian learning for robust matched field processing, *J. acoust. Soc. Am.*, **141**(5), 3411–3420.
- Gerstoft, P. & Tanimoto, T., 2007. A year of microseisms in southern California, *Geophys. Res. Lett.*, 34(20), 2–7.
- Ghosh, A., Vidale, J.E. & Creager, K.C., 2012. Tremor asperities in the transition zone control evolution of slow earthquakes, *J. geophys. Res.*, 117, 1–9.
- Girard, A.J. & Shragge, J., 2020. Automated processing strategies for ambient seismic data, *Geophys. Prospect.*, 68(1), 293–312.
- Gizon, L. & Birch, A.C., 2002. Time-distance helioseismology: the forward problem for random distributed sources, *Astrophys. J.*, 571, 966–986.
- Hanasoge, S.M., 2013. The influence of noise sources on cross-correlation amplitudes, *Geophys. J. Int.*, **192**(1), 295–309.
- Hawthorne, J. & Ampuero, J.-P., 2017. A phase coherence approach to identifying co-located earthquakes and tremor, *Geophys. J. Int.*, 209(2), 623–642.
- Hinich, M.J., 1979. Maximum likelihood estimation of the position of a radiating source in a waveguide, *J. acoust. Soc. Am.*, 66, 480–483.
- Inbal, A., Ampuero, J.P. & Clayton, R.W., 2016. Localized seismic deformation in the upper mantle revealed by dense seismic arrays, *Science*, 354(6308), 88–92.
- Ishii, M., Shearer, P.M., Houston, H. & Vidale, J.E., 2005. Extent, duration and speed of the 2004 Sumatra-Andaman earthquake imaged by the Hi-Net array, *Nature*, 435(7044), 933–936.
- Juretzek, C. & Hadziioannou, C., 2016. Where do ocean microseisms come from? A study of Love-to-Rayleigh wave ratios, J. geophys. Res., 121(9), 6741–6756.
- Kao, H. & Shan, S.J., 2004. The source-scanning algorithm: mapping the distribution of seismic sources in time and space, *Geophys. J. Int.*, 157(2), 589–594.
- Kuperman, W.A. & Turek, G., 1997. Matched field acoustics, *Mech. Syst. Signal Process.*, 11, 141–148.
- Larose, E., Margerin, L., Derode, A., van Tiggelen, B., Campillo, M., Shapiro, N., Paul, A., Stehly, L. & Tanter, M., 2006. Correlation of random wavefields: an interdisciplinary review, *Geophysics*, 71(4), SI11–SI21.
- Lehujeur, M., Vergne, J., Maggi, A. & Schmittbuhl, J., 2017. Ambient noise tomography with non-uniform noise sources and low aperture networks: case study of deep geothermal reservoirs in northern Alsace, France, *Geophys. J. Int.*, 208(1), 193–210.
- Li, L., Tan, J., Schwarz, B., Staněk, F., Poiata, N., Shi, P., Diekmann, L., Eisner, L. & Gajewski, D., 2020. Recent advances and challenges of waveform–based seismic location methods at multiple scales, *Rev. Geophys.*, **58**(1), 1–47.
- Lobkis, O.I. & Weaver, R.L., 2001. On the emergence of the Green's function in the correlations of a diffuse field, *J. acoust. Soc. Am.*, **110**(6), 3011.
- Löer, K., Riahi, N. & Saenger, E.H., 2018. Three-component ambient noise beamforming in the Parkfield area, *Geophys. J. Int.*, 213(3), 1478–1491.
- Ma, Y., Clayton, R.W., Tsai, V.C. & Zhan, Z., 2013. Locating a scatterer in the active volcanic area of Southern Peru from ambient noise crosscorrelation, *Geophys. J. Int.*, **192**(3), 1332–1341.

- Meng, L., Ampuero, J.-P., Stock, J., Duputel, Z., Luo, Y. & Tsai, V.C., 2012. Earthquake in a maze: compressional rupture branching during the 2012 Mw 8.6 Sumatra Earthquake, *Science*, **337**(6095), 724–726.
- Nakata, N. & Beroza, G.C., 2016. Reverse time migration for microseismic sources using the geometric mean as an imaging condition, *Geophysics*, 81(2), KS51–KS60.
- Nakata, N., Snieder, R., Tsuji, T., Larner, K. & Matsuoka, T., 2011. Shear wave imaging from traffic noise using seismic interferometry by crosscoherence, *Geophysics*, **76**(6), 97–106.
- Ozanich, E., Gerstoft, P. & Niu, H., 2020. A feedforward neural network for direction-of-arrival estimation, J. acoust. Soc. Am., 147(3), 2035–2048.
- Retailleau, L., Boué, P., Stehly, L. & Campillo, M., 2017. Locating microseism sources using spurious arrivals in intercontinental noise correlations, *J. geophys. Res.*, **122**(10), 8107–8120.
- Riahi, N. & Gerstoft, P., 2015. The seismic traffic footprint: tracking trains, aircraft, and cars seismically, *Geophys. Res. Lett.*, 42(8), 2674–2681.
- Rost, S. & Thomas, C., 2002. Array seismology: methods and applications, *Rev. Geophys.*, 40(3), 2–1-2-27, .
- Ruigrok, E., Gibbons, S. & Wapenaar, K., 2017. Cross-correlation beamforming, J. Seismol., 21(3), 495–508.
- Sager, K., Ermert, L., Boehm, C. & Fichtner, A., 2018. Towards full waveform ambient noise inversion, *Geophys. J. Int.*, 212(1), 566–590.
- Sager, K., Boehm, C., Ermert, L., Krischer, L. & Fichtner, A., 2020. Globalscale full-waveform ambient noise inversion, *J. geophys. Res.*, 125, 1–17.
- Schmidt, R., 1986. Multiple emitter location and signal parameter estimation, *IEEE Trans. Antenn. Propagat.*, 34(3), 276–280.
- Seats, K.J., Lawrence, J.F. & Prieto, G.A., 2012. Improved ambient noise correlation functions using Welch's method, *Geophys. J. Int.*, 188(2), 513–523.
- Sergeant, A., Małgorzata, C., Lindner, F., Walter, F., Roux, P., Chaput, J., Gimbert, F. & Mordret, A., 2020. On the Green's function emergence from interferometry of seismic wave fields generated in high-melt glaciers: implications for passive imaging and monitoring, *The Cryosphere*, 14(3), 1139–1171.
- Seydoux, L., de Rosny, J. & Shapiro, N.M., 2017. Pre-processing ambient noise cross-correlations with equalizing the covariance matrix eigenspectrum, *Geophys. J. Int.*, 210(3), 1432–1449.
- Shapiro, N.M., Campillo, M., Stehly, L. & Ritzwoller, M.H., 2005. Highresolution surface-wave tomography from ambient seismic noise., *Science* (*New York, N.Y.*), **307**(5715), 1615–8.
- Song, C., Alkhalifah, T. & Wu, Z., 2019. Microseismic event estimation and velocity analysis based on a source-focusing function, *Geophysics*, 84(3), doi:10.1190/geo2018-0205.1.
- Spica, Z.J., Nakata, N., Liu, X., Campman, X., Tang, Z. & Beroza, G.C., 2018. The ambient seismic field at Groningen gas field: an overview from the surface to reservoir depth, *Seismol. Res. Lett.*, 89(4), 1450–1466.
- Stehly, L., Campillo, M. & Shapiro, N.M., 2006. A study of the seismic noise from its long-range correlation properties, *J. geophys. Res.*, **111**(B10), B10306, doi:10.1029/2005JB004237
- Thrastarson, S., van Driel, M., Krischer, L., Boehm, C., Afanasiev, M., van Herwaarden, D.-P. & Fichtner, A., 2020. Accelerating numerical wave propagation by wavefield adapted meshes. Part II: full-waveform inversion, *Geophys. J. Int.*, 221(3), 1591–1604.
- Tromp, J., Luo, Y., Hanasoge, S. & Peter, D., 2010. Noise cross-correlation sensitivity kernels, *Geophys. J. Int.*, 183(2), 791–819.
- Tsai, V.C., 2010. The relationship between noise correlation and the Green's function in the presence of degeneracy and the absence of equipartition, *Geophys. J. Int.*, **182**(3), 1509–1514.
- van Driel, M., Krischer, L., Stähler, S.C., Hosseini, K. & Nissen-Meyer, T., 2015. Instaseis: instant global seismograms based on a broadband waveform database, *Solid Earth*, 6(2), 701–717.
- van Driel, M., Boehm, C., Krischer, L. & Afanasiev, M., 2020. Accelerating numerical wave propagation using wavefield adapted meshes. Part I: forward and adjoint modelling, *Geophys. J. Int.*, 221(3), 1580–1590.

- van Veen, B.D. & Buckley, K.M., 1988. Beamforming: a versatile approach to spatial filtering, *IEEE ASSP Mag.*, **5**, 4–24.
- Walker, K.T. & Shearer, P.M., 2009. Illuminating the near-sonic rupture velocities of the intracontinental Kokoxili Mw 7.8 and Denali fault Mw 7.9 strike-slip earthquakes with global P wave back projection imaging, *J. geophys. Res.*, **114**(2), 1–21.
- Walter, F., Roux, P., Roeoesli, C., Lecointre, A., Kilb, D. & Roux, P.F., 2015. Using glacier seismicity for phase velocity measurements and Green's function retrieval, *Geophys. J. Int.*, 201(3), 1722–1737.
- Wang, Z., Bovik, A.C., Sheikh, H.R. & Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity, *IEEE Trans. Image Process.*, 13(4), 600–612.
- Wapenaar, K., 2004. Retrieving the elastodynamic Green's function of an arbitrary inhomogeneous medium by cross correlation, *Phys. Rev. Lett.*, 93, 254301.
- Wapenaar, K. & Fokkema, J., 2006. Green's function representations for seismic interferometry, *Geophysics*, 71, S133–S146.
- Wassermann, J., 1997. Locating the sources of volcanic explosions and volcanic tremor at Stromboli volcano (Italy) using beamforming on diffraction hyperboloids, *Phys. Earth planet. Inter.*, **104**, 271–281.
- Xu, Z., Mikesell, T.D., Gribler, G. & Mordret, A., 2019. Rayleigh-wave multicomponent cross-correlation-based source strength distribution inversion. Part 1: theory and numerical examples, *Geophys. J. Int.*, 218(3), 1761–1780.
- Yao, H. & van der Hilst, R.D., 2009. Analysis of ambient noise energy distribution and phase velocity bias in ambient noise tomography, with application to SE Tibet, *Geophys. J. Int.*, **179**, 1113–1132.
- Zhang, J. & Yang, X., 2013. Extracting surface wave attenuation from seismic noise using correlation of the coda of correlation, *J. geophys. Res.*, 118(5), 2191–2205.

APPENDIX: ON AMPLITUDES IN MFP

The MFP algorithm presented in the main text only estimates relative delay times between stations, and does not include any term for amplitudes. In contrast, the sensitivity kernels presented in Fig. 6 are constructed using a synthetic forward model which does, by design, model both geometric spreading and an attenuation term. We show here that it is possible to include a relative amplitude term in the MFP algorithm through a more complex replica vector, as has been explored by Corciulo *et al.* (2012) and others.

As described in the main text, the basic MFP algorithm performs a grid search through possible source locations in space. A Δt_{ij} is determined as the difference in traveltimes for a given pair of stations *i* and *j*. We assume a homogeneous velocity model and straight rays, but more sophisticated ray tracing algorithms can be used (Gal & Reading 2018). Corciulo *et al.* (2012) describes an amplitude decay term as well; the distance a wave propagated can be directly mapped to a geometric spreading and attenuation term:

$$A_i = \sqrt{\frac{2v}{\pi \omega r_i}} \exp\left(\frac{-\omega r_i}{2vQ}\right),\tag{A1}$$

where v is the assumed velocity, r_i is the distance from proposed source to a receiver and Q is an attenuation quality factor. The first term represents a decay according to 2-D surface wave geometric spreading and the second term accounts for attenuation. Again, in this MFP framework, we grid search through every location in some spatial domain to test the possibility of a source there. We note that this varies slightly from the equation used by Corciulo *et al.* (2012), but instead directly matches the 2-D analytic Green's functions used for our synthetic modelling (Fichtner *et al.* 2017). This



Figure A1. Panel a shows the stations (triangles) and source (circle) used for synthetic tests. Panel (b) shows the time-shift-only version of the MFP algorithm, panel (c) uses an amplitude correction but also the standard normalization of correlation coefficients in each element of the grid search, and panel (d) uses an amplitude term on the signals with no normalization. Given that the units in each case are different and unimportant, the colourscale is arbitrarily normalized for each.

amplitude term, along with a delta-function time-shift, constitutes the replica vector and might be substituted for what we called W_{ij} in eq. (9):

$$P(x, y) = \sum_{i,j} \int \delta(t - \Delta t_{ij}(x, y)) \,\Delta A_{ij} \,C_{ij}^{obs}(t) \,\mathrm{d}t, \tag{A2}$$

where ΔA_{ij} is the difference in amplitudes computed above, and Δt_{ij} is the difference in expected traveltimes.

This amplitude term, A_i might be used in different ways, as shown in Fig. A1, where we use the same arrangement of 11 stations, but with a single Gaussian source within the domain of the array (panel a). The application of a grid search MFP with only time-shifts, as described in the main text, is shown in panel (b). The first use of an amplitude term might be to correct a signal's amplitude, along with applying time-shifts, but then still normalize the set of correlations by the total observed energy (panel c). This will preserve relative amplitude information, and can be useful in achieving higher resolution while still assuming distant sources are equally likely as nearby sources. A second approach would be to apply the amplitude correction, A_i directly to each unnormalized observation, effectively as if each signal was to be both originating and observed at some location (panel d), which now quickly excludes the possibility of sources further away. In this case, eq. (A2) needs to be further altered to multiply amplitudes for both stations, rather than a define a difference in amplitudes; many MFP works anyways use a formalism of Bartlett processors to apply replica vectors acting on both stations in this way (e.g. Baggeroer *et al.* 1988; Corciulo *et al.* 2012).

With the introduction of both timing and amplitude information, we already see how a full synthetic Green's function might be used to account for an even more complex velocity structure than the one considered here. For example, the term for a delta function might be relaxed to include any number of neighbouring signals from different wave types or different arrivals, reflections, etc. Eq. (A2) still differs from the least-squares misfit defined for the noise-source inversion framework (eq. 10), in that the physics of wave propagation are included in the replica vector instead of a separate C_{ij}^{syn} term. This is perhaps more akin to methods in reverse time migration, that convolve a Green's function to sharpen recovery (e.g. Nakata & Beroza 2016).