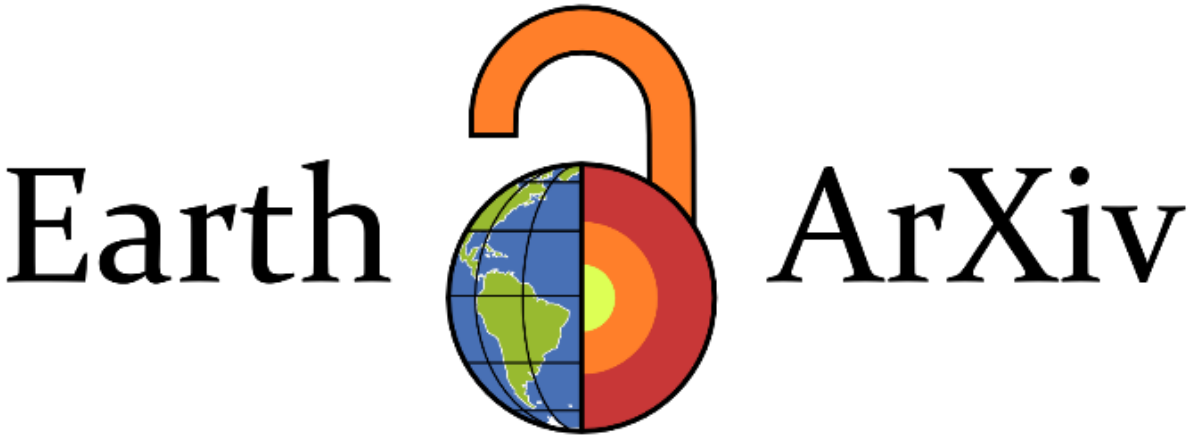


**Title:**
Learning Seismic Wavefield Structure from Regional Arrays Using Self-Supervised Deep Learning

**Authors:**
M. Ronac Giannone* (Southern Methodist University, Dallas, TX, USA; mronacgiannone@smu.edu)
Stephen J. Arrowsmith (Southern Methodist University, Dallas, TX, USA)
Eric Larson (Southern Methodist University, Dallas, TX, USA)

*Corresponding author

**Statement:**
This manuscript is a non-peer-reviewed preprint submitted to EarthArXiv. The manuscript is currently under review at *Seismica*.

# Learning Seismic Wavefield Structure from Regional Arrays with Self-Supervised Deep Learning

**M. Ronac Giannone** * [1], **S. Arrowsmith** [1], **E. Larson** [2]

[1]Department of Earth Science, Southern Methodist University, Dallas, TX, USA, [2]Department of Computer Science, Southern Methodist University, Dallas, TX, USA

Author contributions: *Conceptualization*: M. Ronac Giannone, S. Arrowsmith. *Methodology*: M. Ronac Giannone, E. Larson. *Software*: M. Ronac Giannone. *Validation*: M. Ronac Giannone. *Formal Analysis*: M. Ronac Giannone. *Investigation*: M. Ronac Giannone. *Resources*: S. Arrowsmith. *Writing - Original draft*: M. Ronac Giannone. *Writing - Review & Editing*: M. Ronac Giannone, S. Arrowsmith, E. Larson. *Visualization*: M. Ronac Giannone. *Supervision*: S. Arrowsmith. *Project administration*: S. Arrowsmith. *Funding acquisition*: S. Arrowsmith.

**Abstract** Seismic wavefields recorded across regional arrays exhibit spatial structure governed by propagation physics, including phase alignment, apparent slowness, and source directionality. These properties underpin array processing, yet in sparsely sampled arrays the wavefield is only partially observed, making it difficult to determine how information is distributed across sensors. Particularly, it remains unclear which stations provide the strongest constraints on the wavefield and how spatial coherency can be exploited when inter-station redundancy is limited. In this study, we introduce a U-Net–inspired convolutional-Transformer model that learns an implicit representation of the seismic wavefield from partial array observations through self-supervised reconstruction. We show that the model captures physically meaningful, array-consistent structure, across both body and surface waves as evidenced by improved agreement with ground-truth back azimuth from frequency-wavenumber analysis in challenging reconstruction scenarios. The model further demonstrates robustness to complex waveform conditions, including multi-event windows, temporally offset arrivals, and component-specific transient noise, which it suppresses in favor of coherent wavefield energy. Attention-based analysis reveals that spatially isolated sensors contribute disproportionately useful information for reconstruction, reflecting their unique geometric constraints. These results indicate that our model learns a representation of seismic wavefields with implications for sparse array analysis and future sensor network design.

**Non-technical summary** Seismic sensors are often arranged in arrays to study how energy from earthquakes travels through the Earth. However, when these arrays are sparsely spaced, it can be difficult to understand how each sensor contributes to the overall picture of the wavefield and which sensors provide the most useful information. In this study, we develop a deep learning

*Corresponding author: mronacgiannone@smu.edu

model that learns how seismic waves behave across an array, even when some sensors are missing. Instead of simply filling in missing data, the model captures the underlying physical structure of how waves move and interact. We show that the model can accurately recover important properties such as the direction from which waves arrive, even in difficult cases where the data are incomplete or noisy. It also performs well in complex situations, such as when multiple events occur in the same time window or when signals are shifted in time, and it can ignore short bursts of noise that do not reflect true wave behavior. Our results also reveal that more isolated sensors often provide especially valuable information. These findings suggest that the model can help improve how seismic arrays are used and designed in the future.

# 1 Introduction

Seismic wavefields recorded across spatially distributed sensors exhibit a property known as coherency, which describes how wavefield measurements vary across space due to propagation physics. Coherency underpins array seismology, governing observable quantities such as apparent slowness, back azimuth, and phase alignment across sensors. However, coherency is extremely difficult to model in practice, as it is strongly influenced by scattering from complex three-dimensional heterogeneity at sub-wavelength scales (Darrh et al., 2019). As a result, it is typically represented using simplified empirical models that remain poorly constrained by observational data. This limitation raises a fundamental question: how is information about the seismic wavefield distributed across an array, and which sensors provide the strongest constraints on its spatial structure?

One approach to probing wavefield coherency is through seismic data reconstruction, where missing observations are inferred from available measurements typically in dense acquisition settings, such as shot-gather interpolation and missing-data recovery in nodal networks. Several studies have addressed reconstruction using traditional signal processing approaches, including wavelet-curvelet compressive sensing and localized rank reduction (Chen et al., 2019; Muir and Zhan, 2021). More recent work has applied deep learning techniques to improve reconstruction performance. A key consideration in these approaches is the type of contextual information used to reconstruct missing waveforms. For example, Gaharwar et al. (2023) introduced a neural network that reconstructs waveform gaps in time-series data using temporal and spectral context, whereas other studies address shot-gather interpolation and missing trace reconstruction using the spatial context provided by neighboring traces (Wang et al., 2019; He et al., 2022; Yu and Wu, 2022).

Self-supervised learning has been a common strategy in this setting, where models are trained by masking observed data and learn to recover missing portions from available waveform segments (Abedi and Pardo, 2022; Huang et al., 2022). Consequently, these models learn the structure of seismic wavefields from partial observations. In dense networks, effective reconstruction leverages strong spatiotemporal coherency and substantial spatial redundancy, which can also improve downstream tasks such as event location (Meng et al., 2022; Li and Schwarz, 2024). In contrast, regional seismic arrays are much more sparsely distributed, so the available context is weaker and the reconstruction problem becomes more directly constrained by the underlying wavefield physics. The question is

not only how to fill gaps, but what the array wavefield itself reveals about phase alignment, apparent slowness, and source directionality when redundancy is limited. This perspective is consistent with wavefield-level reconstruction methods that treat the array as a coherent structure rather than a collection of independent traces (Abedi et al., 2024).

Regional seismic arrays, with typical apertures of 1-2 km, consist of sensors deployed in fixed geometries to detect signals from events at distances of tens to hundreds of kilometers (Gibbons et al., 2010; Stump et al., 2022). Sensor spacing is typically designed relative to the wavelengths of the target signals expected at these ranges. As a result, these arrays often include only a few to tens of sensors within a given aperture, making them significantly sparser than dense seismic networks. Dense networks prioritize high spatial sampling to enable subsurface imaging and local characterization, whereas regional seismic arrays, with larger apertures and sparser spacing, are optimized for detecting and characterizing distant sources using array processing (Rost and Thomas, 2002). These types of deployments are important because they have been used extensively on a global scale as part of the International Monitoring System (IMS) to detect and locate nuclear explosions for international treaty verification (Christie and Campus, 2010; Koper, 2020). These sparse and dense settings therefore differ fundamentally in how information about the wavefield is distributed across sensors and in the extent to which spatial coherency constrains the underlying wavefield. This sparsity makes regional arrays a particularly informative setting for studying wavefield coherency, as limited spatial redundancy forces the underlying propagation physics to govern how information is distributed across sensors (Shi et al., 2021).

Recent work has begun to probe wavefield structure in regional arrays using data-driven approaches. One study employed an encoder-decoder network to learn transfer functions between pairs of nearby sensors, mapping waveforms from a fixed reference station to a target station at separations of approximately 200 m (Klinge et al., 2025). A follow-up study introduced the concept of Virtual Seismic Arrays, which predict full array recordings from a single reference station and can provide data for stations that are temporarily offline or no longer deployed (Klinge et al., 2026). While these approaches demonstrate that aspects of the wavefield can be inferred from partial observations, they remain fundamentally reference-centered. Consequently, reconstruction quality depends on the availability and representativeness of a single station, and the full spatial structure of the wavefield across the array is not explicitly leveraged. A key question is whether such models learn a physically meaningful representation of wavefield coherency across the array, or instead approximate a collection of pairwise transfer relationships between sensors.

In this study, we extend this approach by developing a conditional deep learning model designed to learn an implicit representation of regional seismic wavefields from partial array observations. Instead of relying on a fixed reference sensor or learning pairwise transfer functions, the model reconstructs missing three-component waveforms by exploiting spatial coherency across the entire array. The architecture combines component-wise encoder-decoder branches with Transformer blocks that capture cross-sensor and cross-component dependencies explicitly conditioned on inter-station geometry. Trained on 30,000 local earthquakes recorded across a 10-sensor, 1-km aperture array, the model reconstructs full 60-second, three-component recordings that include both body and surface waves. This allows the network to internalize the underlying structure of the propagating wavefield itself. Consequently, reconstruction is treated as a mechanism for learning and interrogating the spatial structure of the seismic wavefield.

Reconstruction quality is first evaluated using normalized cross-correlation between observed and reconstructed waveforms, exceeding 0.98 at well-connected sensors and remaining above 0.80 at the most isolated sensor. More importantly, even in these challenging scenarios, reconstructed waveforms preserve physically meaningful wavefield properties. In particular, frequency-wavenumber (*fk*) analysis applied to reconstructed data yields closer agreement with ground-truth back azimuth and increased array coherence for both *P*- and *S*-wave windows, as well as their associated coda. The model further demonstrates robustness under complex waveform conditions, including multi-event windows and temporally offset arrivals, and selectively suppresses component-specific transient noise in favor of spatially coherent wavefield energy. Attention-based analysis and sensor-removal experiments reveal that spatially isolated stations, although hardest to reconstruct, impose the strongest constraints on the wavefield due to their unique geometric positioning. Together, these results indicate that the model learns a physically meaningful, array-consistent representation of the seismic wavefield, capturing its spatial organization rather than acting as a purely interpolative mapping, with implications for both sparse array analysis and sensor network design.

# 2 Data

## 2.1 Acquisition

Seismic waveform data were obtained from the Piñon Flat Observatory (PFO) array in southern California via the International Federation of Digital Seismograph Networks (FDSN) web services. An accompanying earthquake catalog was constructed using event metadata from the Southern California Earthquake Data Center (SCEDC). These combined datasets provide three-component waveform recordings and corresponding event parameters used for model training and evaluation. Data span local earthquakes occurring between August 2014 and December 2025 within 100 km of the PFO array, with only events of magnitude $\geq 1$ considered. Furthermore, for each cataloged event, waveform windows were centered relative to the predicted *P*-wave onset at the array. Predicted onset times were computed using the *iasp91 TauP* travel-time model from event origin time, hypocentral depth, and epicentral distance between the event and the array reference location. The earliest arrival among *P*-phases was used as the reference onset. The final waveform window extended from 10 seconds before to 50 seconds after this predicted *P* arrival.

To ensure a consistent regional array geometry across the catalog, we first evaluated waveform availability for all 13 PFO stations (BPH01 – BPH13) across the three required components (BHZ, BHN, and BHE). For each event, waveform recordings were scanned to determine which stations contained all three components, and a station was counted as missing for that event if one or more required components were absent. We then tabulated the number and percentage of events for which each station was incomplete. Stations missing one or more components in at least 10% of cataloged events were treated as persistently unreliable and removed from the final database. This screening identified stations BPH04, BPH05, and BPH12 as chronically incomplete, and they were excluded from all subsequent processing, reducing the working array from 13 to 10 sensors. The resultant array is displayed in Figure 1.

After removing these persistently incomplete stations, each event was subjected to additional quality-control checks to ensure consistent array dimensionality and waveform length. For each component, the retained traces were required to include exactly 10 unique stations and have a fixed length of 2400 samples, corresponding to 60-second windows at a 40 Hz sampling rate. Events were further rejected if any trace contained more than 50% zero-

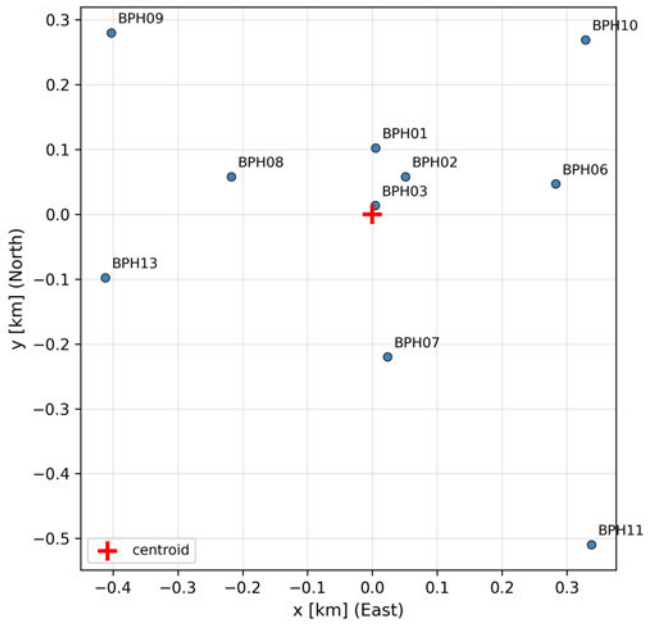


**Figure 1** Geometry of the PFO seismic array after removal of stations with persistent data gaps. Sensor locations are shown in a local Cartesian coordinate system (east-north, km) relative to the array centroid (red cross). The resulting 10-station configuration spans a range of inter-station distances, providing a multi-scale spatial sampling of the wavefield. This diversity in sensor spacing enables evaluation of model performance across both closely spaced and more isolated stations.

amplitude samples, indicating potential data corruption or invalid waveform content. Event identifiers failing any of these criteria were added to an automated removal list and excluded from the final training dataset. Finally, we manually reviewed each event and applied array processing. Events were flagged and subsequently removed if their traces contained impulsive noise or if estimated back azimuths did not align with the ground-truth source location. In total, we compile a dataset consisting of 30,000 local earthquakes.

## 2.2 Spatial Coherency

Since the ability to reconstruct missing sensors depends fundamentally on the degree to which waveforms remain coherent across the array, we first quantify the spatial coherency of the wavefield. To characterize this coherency, we analyzed the full 60-second waveform windows in the BHZ component over the 0.5 – 10 Hz band. The lower bound of 0.5 Hz was selected to suppress long-period microseismic energy, while the upper bound of 10 Hz was chosen because higher frequencies showed a marked increase in impulsive noise contamination within the recordings. For each event, traces were first amplitude-normalized on a station-by-station basis, and coherency was quantified using the maximum normalized cross-correlation between pairs of sensors over the full event window. An event-level mean coherency was then computed by averaging the maximum pairwise cross-correlations across all station pairs. These values were used to examine how spatial coherency varies with source azimuth, epicentral distance, and event magnitude in Figures 2a, b. Figure 2a suggests that the wavefield is not isotropic insofar that coherency depends strongly on direction. Moreover, Figure 2b shows that coherency is largely independent of source size, reinforcing that it is governed by propagation, not amplitude.

To evaluate the frequency dependence on spatial coherency, the waveforms were further divided into four band-pass intervals: 0.5 – 1, 1 – 2, 2 – 5, and 5 – 10 Hz. For each band, event-level mean coherency was recomputed and

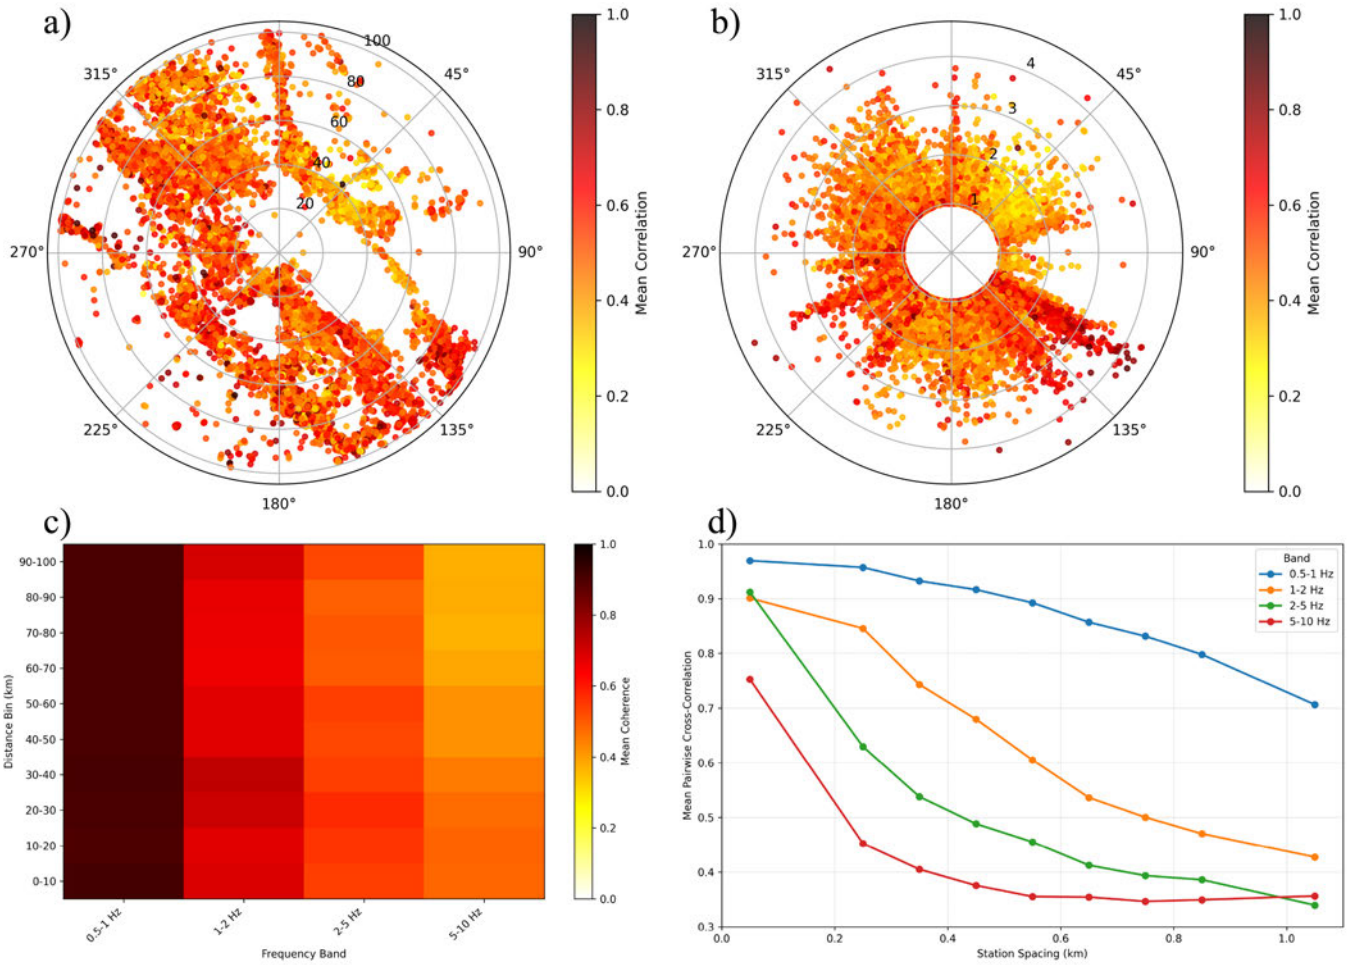


**Figure 2** Spatial coherency of the PFO array wavefield in the BHZ component over the 0.5 – 10 Hz band. a) Event-level mean coherency as a function of source azimuth and epicentral distance. b) Event-level mean coherency as a function of source azimuth and event magnitude. c) Mean coherency as a function of source distance and frequency band. d) Mean pairwise maximum cross-correlation as a function of inter-station spacing for the same frequency bands.

then aggregated as a function of source distance to form the frequency-versus-distance heat map in Figure 2c. Additionally, station geometry from the array site file was used to calculate inter-station spacing for every sensor pair, and mean pairwise maximum cross-correlation was summarized within spacing bins for each frequency band. Results are displayed in Figure 2d. All in all, these plots demonstrate that spatial coherency at PFO is governed by wavelength-dependent propagation physics and array geometry, resulting in a heterogeneous distribution of information across sensors, with secondary dependence on azimuth and weak dependence on magnitude. Based on these trends, subsequent analyses focus on a 0.5 – 5 Hz bandpass, which captures the range of frequencies where spatial coherency remains sufficiently strong to constrain the wavefield while avoiding the rapid decorrelation observed at higher frequencies. Although the bulk of the analysis is conducted within this band, we also report selected results up to 10 Hz to illustrate performance in more weakly coherent regimes.

# 3 Methods

## 3.1 Preprocessing

Prior to model training, all waveform records were preprocessed consistently across events and sensors. Each trace was first cosine tapered, then demeaned and linearly detrended to reduce edge effects, constant offsets, and long-

period linear drift. The data was subsequently bandpass filtered between 0.5 and 5 Hz using a fourth-order zero-phase Butterworth filter applied along the time axis. The preprocessed event database was partitioned into training and testing sets using a random 80/20 split at the event level, ensuring that no individual event appears in both partitions. Be that as it may, due to the spatial and temporal clustering of seismicity, events in the training and testing sets may occur close to one another in space and time, resulting in overlapping distributions. The spatiotemporal distribution of training and testing events is shown in Figure S1, illustrating that while events do not overlap between partitions, they do occupy similar regions of seismicity. For the purposes of this study, which focuses on evaluating whether the model learns physically meaningful representations of wavefield structure within a fixed array, this in-domain split is appropriate. Assessment of generalization to more distinct source distributions, including new regions and event types, is deferred to future work that explicitly investigates transferability in a deployment-oriented setting. Our dataset splitting approach was performed with a fixed random seed to ensure reproducibility, and the associated event metadata were subset in parallel so that waveform records and catalog attributes remained aligned in both partitions.

To address the large variability in raw amplitudes within and across events, waveforms were standardized using robust global scaling parameters computed separately for each component from the training set. Specifically, component-wise centers and scales were estimated using the median and median absolute deviation (*MAD*). These parameters were used in a global inverse-hyperbolic-sine (*asinh*) scaling transform that preserves waveform sign while compressing large amplitudes and reducing sensitivity to extreme values. This normalization preserves physically meaningful amplitude relationships while compressing the dynamic range of seismic signals, encouraging the model to learn spatial structure of the wavefield rather than dataset-specific amplitude variability. Training examples were regenerated at each epoch using a stateless data pipeline, meaning that all random masking and subsequent augmentation operations were driven by explicit seeds rather than by an internal random state. This allows each epoch to represent a new masking realization while preserving full reproducibility of the training process.

At the epoch level, events were resampled with magnitude-based oversampling so that earthquakes above a specified magnitude threshold (mag $\geq 2$) constituted a fixed fraction of the training set (now 25%, originally representing 10% of all events), preventing the model from being dominated by the more numerous smaller events. This approach improves amplitude reconstruction, mitigating the tendency of the model to under-predict higher-amplitude signals. At the sample level, waveforms were augmented with random time shifts ($\pm 5$ seconds), after which a single sensor was randomly selected and masked across all three components for the full time window. The selected sensor was zeroed out across all components and time samples, while the remaining observed traces were transformed into the globally scaled model space and perturbed with low-amplitude additive noise (2.5%). The resulting waveform tensor was concatenated with a binary observation mask, allowing each component branch to distinguish between observed and missing data. Prediction targets were constructed separately for each component by pairing the fully scaled waveform with the associated masked-sensor indicator. Validation examples were generated using the same global component-wise scaling and masking formulation as the training data, but without stochastic augmentation. In particular, no additive noise, time shifting, or oversampling was applied to the validation set, ensuring that model performance was evaluated on a fixed, independent, and reproducible sensor-masking configuration.

## 3.2 Model Inputs

We provide the model with a matrix of inter-station distances that encodes the fixed geometry of the array. This distance matrix is used within each Transformer block as a learned spatial bias added directly to the attention scores before the *softmax* operation. In this way, attention between pairs of sensor-component tokens is conditioned not only on their learned waveform representations, but also on their physical separation within the array. This geometric conditioning is motivated by the fact that waveform similarity and redundancy within a regional array are not spatially uniform. Inter-station spacing influences the degree to which one sensor can constrain another through the spatial coherency of the wavefield. Incorporating distance explicitly allows the model to learn cross-sensor dependencies in a manner that is informed by the physical layout of the array rather than inferred from waveform similarity alone, providing a direct inductive bias toward learning spatial coherency.

Figure 3a shows an example three-component waveform record across all 10 sensors, yielding 30 sensor-component traces. Following, Figure 3b displays the corresponding inter-station distance matrix with an aperture $> 1$ km, apparent between sensors BPH09 and BPH11. Along with the observed mask, these representations form the inputs to the proposed neural network introduced in the following subsection. The model input consists of three-component waveform data concatenated with corresponding binary mask channels, resulting in a six-channel input tensor of shape $[S, T, 6]$, where the mask explicitly indicates which samples are observed and which are withheld. During training and evaluation, a single sensor is masked per example by removing all three components across the full time window, and the model is tasked with reconstructing the corresponding waveform using the remaining array observations. These inputs provide both waveform observations and explicit geometric context, allowing the model to learn how seismic wavefields vary across space rather than relying solely on local or pairwise relationships.

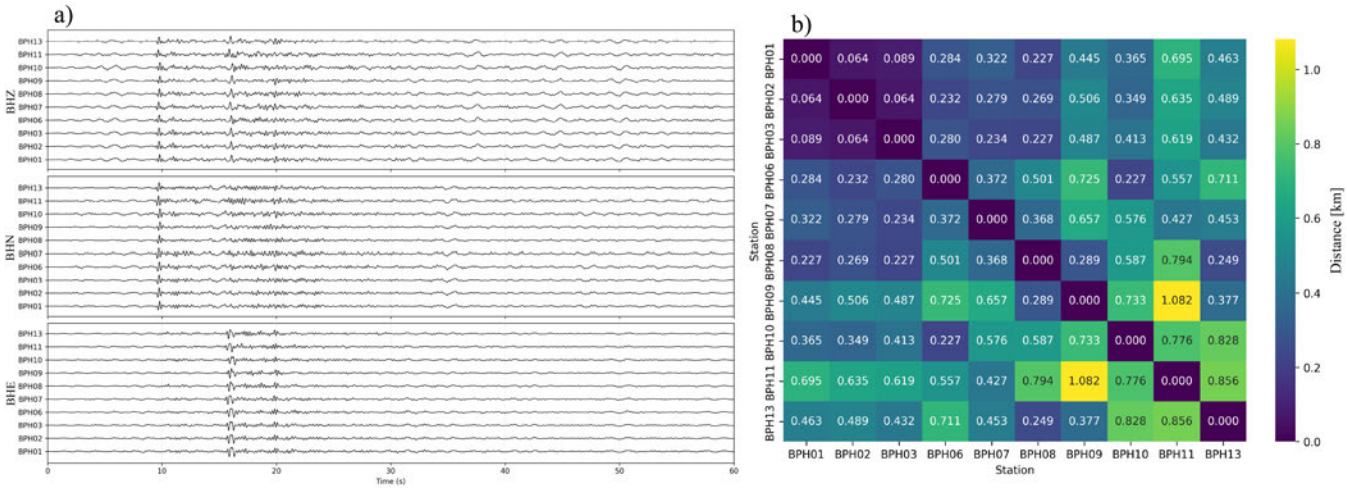


**Figure 3** Inputs to the model. Panel a) plots a three-component waveform record across the PFO array produced by a local earthquake. Panel b) shows the inter-station distance matrix for the array.

## 3.3 Neural Network Architecture

Our proposed neural network architecture, referred to as the Virtual Array Synthesis Architecture (*VASA*), is designed to learn an implicit representation of seismic wavefield coherency across regional arrays through self-supervised masked reconstruction (Figure 4). The architecture combines convolutional feature extraction layers, including pre-activation-style residual blocks, with Transformer blocks that model long-range spatial dependencies across sen-

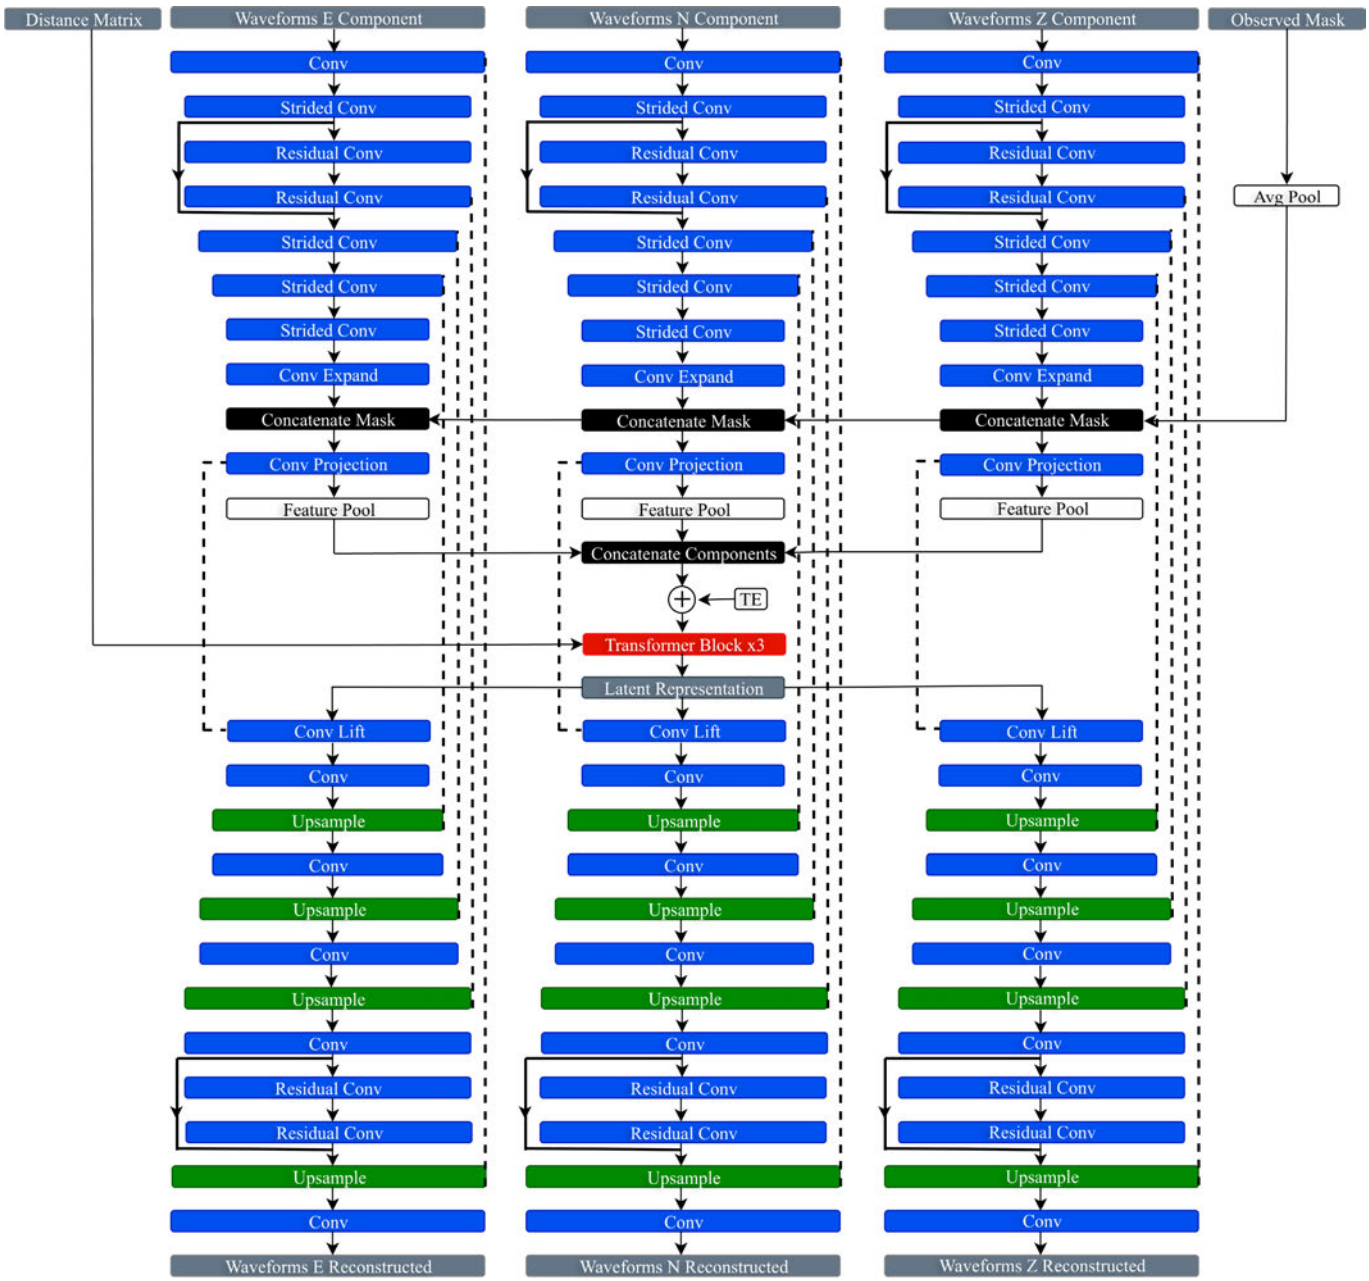


**Figure 4** High level representation of the Virtual Array Synthesis Architecture (*VASA*). Blue layers represent convolutional operations, green represent upsampling layers using bilinear interpolation, and red denotes the Transformer processing blocks. Dashed lines represent skip connections between the encoder and decoder, while grey layers denote inputs, outputs, and fused latent space. Feature pooling and average pooling layers are displayed in white. *TE* represents the joint token embeddings.

sors within the array. The model processes each seismic component using separate encoder-decoder branches, preserving component-specific waveform structure prior to cross-component fusion. This design preserves physically meaningful features related to the source and propagation path, such as polarization, arrival characteristics, and amplitude. By avoiding early fusion across components, the model is not forced to learn a shared local representation prematurely. The architectural choice is also supported empirically. Preliminary experiments in which all three components were processed jointly through a shared encoder-decoder branch resulted in substantially worse reconstruction performance compared to the component-wise multi-branch design. We attribute this degradation to premature feature fusion, in which component-specific information is combined too early and thus blurred, limiting the model's ability to preserve and exploit physically meaningful distinctions in the wavefield. This design encour-

ages the model to learn a global, array-consistent representation of the wavefield, rather than approximating pairwise transfer relationships between individual sensors.

Within each component branch, waveform features are extracted using an encoder that progressively reduces the temporal dimension using convolutional strides while increasing feature depth. Early layers use larger convolutional kernels to capture low-frequency, long-duration signal characteristics, while later layers use smaller kernels to resolve higher-frequency structure after temporal downsampling. This multi-scale design enables the network to capture physically meaningful waveform features, from localized arrivals to broader event-level propagation patterns. The temporal downsampling operations are followed by a $1\times1$ convolutional expansion layer that increases the representational capacity of each component branch (i.e., increased number of filters x4). This operation projects local waveform features into a higher-dimensional embedding space, enabling richer feature representations for subsequent interaction and aggregation layers. Following, the observed mask input is downsampled through average pooling to align with the reduced temporal resolution of the convolutional feature map and concatenated with the expanded features. This combined representation is then passed through a $1 \times 1$ convolutional projection layer that mixes features at each time step, allowing the model to explicitly condition feature interactions on observed and missing waveform regions during feature interaction.

To interface the convolutional encoder with the Transformer blocks, we apply a learned feature pooling layer that converts each component-specific feature map into a set of sensor tokens. The encoder output has shape [$S$, $T'$, $F$], where $S$ is the number of sensors (10), $T'$ the downsampled temporal dimension (150), and $F$ the feature dimension (32). To satisfy the Transformer's input requirements, the feature dimension is collapsed using a learned, time-varying attention pooling operation, yielding a representation of shape [$S$, $T'$]. Specifically, a convolutional layer produces intermediate activation scores that are projected and normalized with a *softmax* function along the feature dimension, generating attention weights that define the relative importance of each feature channel at every time step. The pooled representation is then computed as a weighted sum of the feature map along the feature axis, allowing feature-dependent information to contribute through learned weights. This operation preserves salient waveform characteristics while reducing dimensionality, and thus retains the downsampled temporal structure as the embedding dimension of each sensor token. This formulation is analogous to the attention-based pooling used in previous waveform reconstruction studies that achieved high reconstruction fidelity, where the frequency axis is collapsed for Transformer processing while preserving spectro-temporal structure (Ronac Giannone et al., 2026). This representation treats each sensor as a compact summary of its local wavefield observation, enabling subsequent modeling of spatial relationships across the array.

The pooled representations from each component branch are then concatenated along the token axis to form a unified token set of shape [$3S$, $T'$], where $S$ is the number of sensors and $T'$ the reduced temporal dimension. For a 10-sensor array, this produces 30 sensor-component tokens corresponding to the *E*, *N*, and *Z* components at each station. This ordering groups tokens by component across all sensors, enabling joint modeling of cross-component and cross-sensor interactions. To provide the model with explicit information about sensor identity and component type, learned station and component embeddings are added to the token representation. The station embeddings encode which array element each token belongs to, while the component embeddings distinguish whether the token

corresponds to the vertical, north, or east channel. These token embeddings are broadcast across the batch dimension and added directly to the pooled latent representation, allowing the Transformer to condition attention on both sensor identity and component-specific wavefield structure.

The concatenated token sequence is processed by a stack of spatially aware Transformer blocks that model long-range dependencies across sensors and components. Each block follows a standard Transformer architecture consisting of multi-head self-attention, a position-wise feedforward network, residual connections, and layer normalization. To incorporate array geometry directly into the attention mechanism, we introduce a spatial bias derived from the inter-station distance matrix. Specifically, pairwise distances between sensors are mapped through a small multilayer perceptron to produce a learned bias term, which is added to the attention scores prior to the *softmax* operation. This allows attention weights to be conditioned not only on feature similarity but also on physical separation between sensors, providing a direct inductive bias toward learning spatial coherency in the wavefield. As a result, the Transformer learns interactions that reflect the spatial coherency of the seismic wavefield, giving the model the capacity to capture directionality, phase relationships, and propagation structure across the array.

After the Transformer blocks, the updated token sequence is separated back into distinct component representations and passed to component-specific decoder branches. Each decoder first reshapes the latent representation from the Transformer output into a sensor-time feature map and applies a $1 \times 1$ convolutional lifting layer to enrich the representation and align it with the feature dimensionality of the projected skip connections. The temporal dimension is then progressively restored to its original size using bilinear upsampling. *VASA* follows a U-Net-inspired encoder-decoder design with skip connections between corresponding downsampling and upsampling stages (Ronneberger et al., 2015). This design is particularly effective for reconstructing high-frequency waveform content. At each scale, the corresponding encoder skip representation is first projected through a $1 \times 1$ convolution to align its feature dimension with the decoder pathway, and is then concatenated with the upsampled decoder representation. This preserves fine-scale waveform information that may otherwise be lost during temporal compression while allowing the decoder to integrate multi-resolution context in a consistent feature space. An additional residual convolutional block is applied to further stabilize intermediate reconstruction before the final upsampling stage. A final linear convolution produces reconstructed waveform segments for each component.

*VASA* is trained over 1000 epochs using separate masked mean-squared-error (MSE) losses for each component output. During training, a single sensor is masked per example, and the model is tasked with reconstructing the corresponding waveform using the remaining array observations. Each target contains both the fully scaled waveform and a binary mask identifying the masked sensors, so that the loss is evaluated only over samples belonging to the missing sensor waveforms. This focuses optimization on reconstructing unobserved data rather than reproducing already observed waveform segments. During training and validation, reconstruction quality is further monitored using masked root-mean-square error (RMSE) and masked zero-lag Pearson correlation, each computed only over the masked sensor waveforms.

## 3.4 Model Evaluation

We evaluate reconstruction quality using both zero-lag Pearson correlation and normalized waveform cross-correlation. The former measures the linear similarity between observed and reconstructed waveforms without allowing for any

temporal shift, and is computed only over the masked sensor samples. As such, it provides a strict measure of how accurately the model recovers waveform amplitudes and phase alignment at each time step. High zero-lag correlation therefore indicates that the reconstructed waveform matches both the timing and shape of the true signal. In contrast, normalized waveform cross-correlation evaluates similarity as a function of time lag by computing the correlation between normalized signals over all possible temporal offsets (Aki and Richards, 2002). The maximum value of this function provides a measure of waveform similarity that is invariant to small timing errors, allowing us to assess whether the reconstructed waveform captures the correct signal structure even if it is slightly misaligned in time. These metrics provide complementary perspectives on reconstruction performance. Zero-lag correlation emphasizes exact waveform recovery, while cross-correlation captures broader structural similarity and phase alignment. Differences between the two metrics therefore provide insight into whether reconstruction errors arise from amplitude distortions or temporal misalignment, and more generally reflect how well the model has learned physically meaningful wavefield characteristics such as arrival timing, phase alignment, and waveform morphology. All waveform-based metrics are computed under a single-sensor masking framework, where each sensor is withheld individually and reconstructed using the remaining array observations.

In addition to waveform-based metrics, we employ traditional array processing techniques to further evaluate whether *VASA* preserves physically meaningful wavefield structure. Array processing enables beamforming methods that exploit differential travel times between sensors to coherently align and stack signals across the array (Rost and Thomas, 2002, 2009). In particular, *fk* analysis estimates the full slowness vector by performing a grid search over trial slownesses in the frequency domain, providing measures of signal directionality, trace velocity, and coherence. We apply *fk* analysis to both observed and reconstructed waveforms, focusing on the *P*- and *S*-wave phases recorded at the most spatially isolated sensor. This setting represents the most challenging reconstruction scenario and therefore provides a stringent test of whether the model preserves physically consistent wavefield structure. We compare back azimuth and trace velocity estimates derived from reconstructed waveforms to those obtained from the observed data and to ground-truth source locations provided by the SCEDC. Agreement, or improvement, between these estimates indicates that the reconstructed waveforms, despite exhibiting lower waveform-based metrics, still retain the directional and kinematic structure of the underlying seismic wavefield. This provides a physically grounded validation that the model has learned a coherent representation of the seismic wavefield rather than acting as a purely interpolative mapping between sensors.

As a final evaluation step, we analyze the learned attention patterns within the Transformer blocks to assess the relative importance of individual sensors during reconstruction. To do so, we generate all possible single-sensor masking scenarios for the validation dataset, such that each event is replicated $S$ times with a different sensor withheld in each instance. The trained model is then evaluated on this expanded dataset, and cross-component attention weights are extracted from the Transformer blocks, focusing specifically on interactions associated with the masked sensor. The resulting attention maps, defined at the sensor-component token level, are subsequently aggregated to produce a sensor-level attention matrix by averaging across components. This provides a proxy for how strongly each sensor contributes to reconstructing other sensors in the array. Row-normalized attention weights are used to interpret the relative influence of contributing sensors for a given masked sensor, allowing for the direct compari-

son of their importance. This analysis provides insight into how the model distributes information across the array and reveals which sensors contribute the most to reconstruction under different masking scenarios. In particular, it allows us to identify whether spatially isolated stations, which are more difficult to reconstruct, nonetheless play a disproportionate role in informing the reconstruction of other sensors due to their unique geometric constraints, thereby providing a physically interpretable measure of how information about the seismic wavefield is distributed across the array. This interpretation is further validated through targeted sensor-removal experiments, which directly quantify the impact of individual sensors on reconstruction performance.

# 4 Results

## 4.1 Training

We present masked single-sensor *VASA* training results in Figure S2. Overall, training converged smoothly across all three components, with masked zero-lag validation correlations increasing rapidly during the first $\sim$100 epochs and then gradually approaching stable values near 0.90. Corresponding masked RMSE decreased consistently throughout training, with the *Z*-component exhibiting the lowest RMSE and the *E*-component the highest. Training and validation losses followed similar trajectories, indicating stable optimization without strong evidence of overfitting. The best model was selected at epoch 993 based on the maximum summed validation correlation across all components, which reached 2.691, averaging 0.897 per component.

Normalized waveform cross-correlation is implemented to evaluate reconstruction performance separately for each component using the maximum cross-correlation, best lag, root-mean-square (RMS) ratio, and peak-amplitude ratio (Figure S3). For each event, every sensor is masked individually, the corresponding reconstruction metrics are computed at the masked sensor, and the results subsequently averaged across sensors for each component. The distributions of maximum cross-correlation values are tightly clustered around 0.91, indicating consistently high waveform similarity across the test set with minimal difference between components. The distribution of best lag values is sharply peaked near zero for all components, demonstrating that reconstructed waveforms are well aligned in time and preserve arrival structure. Amplitude recovery, which previous studies such as Klinge et al. (2025, 2026) and Gaharwar et al. (2023) note as particularly challenging, is also strong. Peak amplitude ratios are close to unity, with means of 1.052, 1.046, and 1.064 for *Z*, *N*, and *E* components, respectively, indicating that the model accurately recovers dominant waveform peaks. The RMS ratios, which reflect amplitude agreement across the entire 60-second waveform, are slightly lower, with means of 0.929, 0.925, and 0.921, suggesting that the model modestly underestimates the average waveform energy. All in all, these results suggest that *VASA* captures both the shape, timing, and amplitude of seismic waveforms with high accuracy, and that reconstruction performance is consistent across components despite differences in polarization and wavefield characteristics.

For each station, reconstruction quality was evaluated by computing the maximum normalized cross-correlation between observed and reconstructed waveforms across all events and components. Across individual sensors, performance shows a clear dependence on array geometry and inter-station separation. Cross-correlation remains high for less isolated sensors (generally $>$0.90, with a max around 0.98 at BPH01/02/03), but decreases for more isolated sensors, with the lowest values near $\sim$0.80 at BPH11 (Figure S4a). A similar trend is observed in the RMS ratio, where

sensors with dense local coverage reach values near 1.0, while more isolated stations exhibit reduced values ranging from ∼0.75 to 0.85 (Figure S4b). In contrast, peak amplitude ratios remain close to one for nearly all sensors, even for those with lower correlation and RMS agreement, indicating that *VASA* consistently recovers the dominant waveform peaks (Figure S4c). These results suggest that reconstruction errors at challenging sensors primarily arise from their spatial isolation and reduced recovery of distributed waveform energy, highlighting the role of spatial redundancy in constraining full waveform structure. For visualization, we present the observed and reconstructed waveform per component for a local earthquake recorded by BPH02 (Figure 5). We also show an example where the model reconstructs both body *and* prominent surface wave structure with high fidelity (Figure S5).

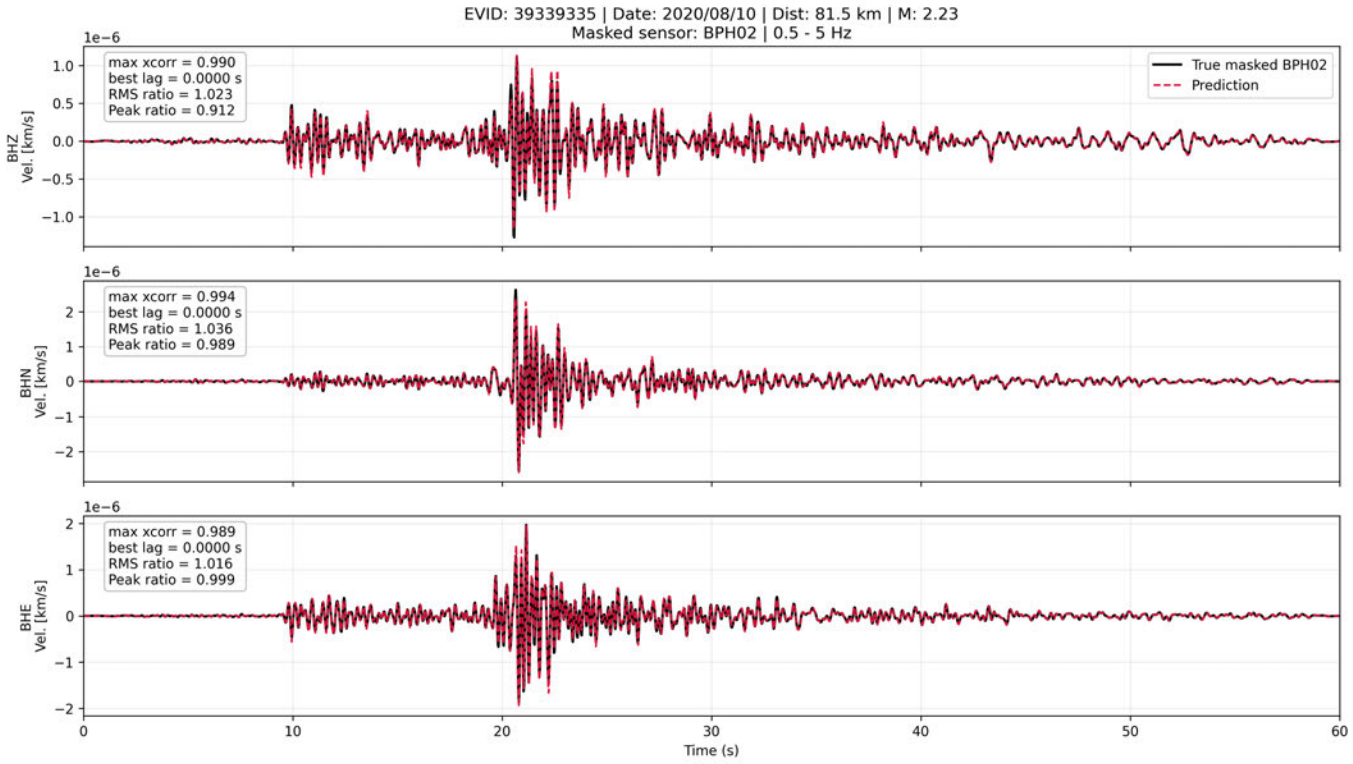


**Figure 5** Observed and reconstructed three-component waveforms for a representative event with BPH02 masked during inference. Black lines show the true waveform at the masked sensor, while red dashed lines indicate the model reconstruction. Reconstruction quality is quantified using maximum normalized cross-correlation, best lag, RMS amplitude ratio, and peak amplitude ratio for each component. The model accurately recovers waveform shape, timing, and amplitude across all components, with cross-correlation values ∼0.99 and negligible time lag

## 4.2 Array Processing Validation

While waveform-based metrics provide a quantitative measure of reconstruction accuracy, they do not directly assess whether the reconstructed signals preserve the physically meaningful structure of the seismic wavefield. To address this, we use *Cardinal*, an array processing software, to evaluate the extent to which spatial coherency and propagation characteristics are retained in the reconstructed test dataset (Ronac Giannone et al., 2025). Specifically, we compare back azimuth estimates derived from *fk* analysis of observed and reconstructed waveforms, using both *P*- and *S*-wave windows, and benchmark these results against ground-truth back azimuths obtained from event locations. We focus this analysis on station BPH11, which exhibits the lowest reconstruction performance in waveform-based metrics and therefore represents the most challenging reconstruction scenario. This provides a stringent test of whether the reconstructed waveforms contain physically meaningful information or simply reflect interpolation of neighboring

signals. Agreement between observed and reconstructed back azimuth estimates, particularly in this worst-case setting, indicates that the model preserves the directional and kinematic structure of the underlying seismic wavefield. More broadly, this analysis provides evidence that *VASA* learns a coherent, physically consistent representation of the seismic wavefield, rather than merely reproducing waveform amplitudes.

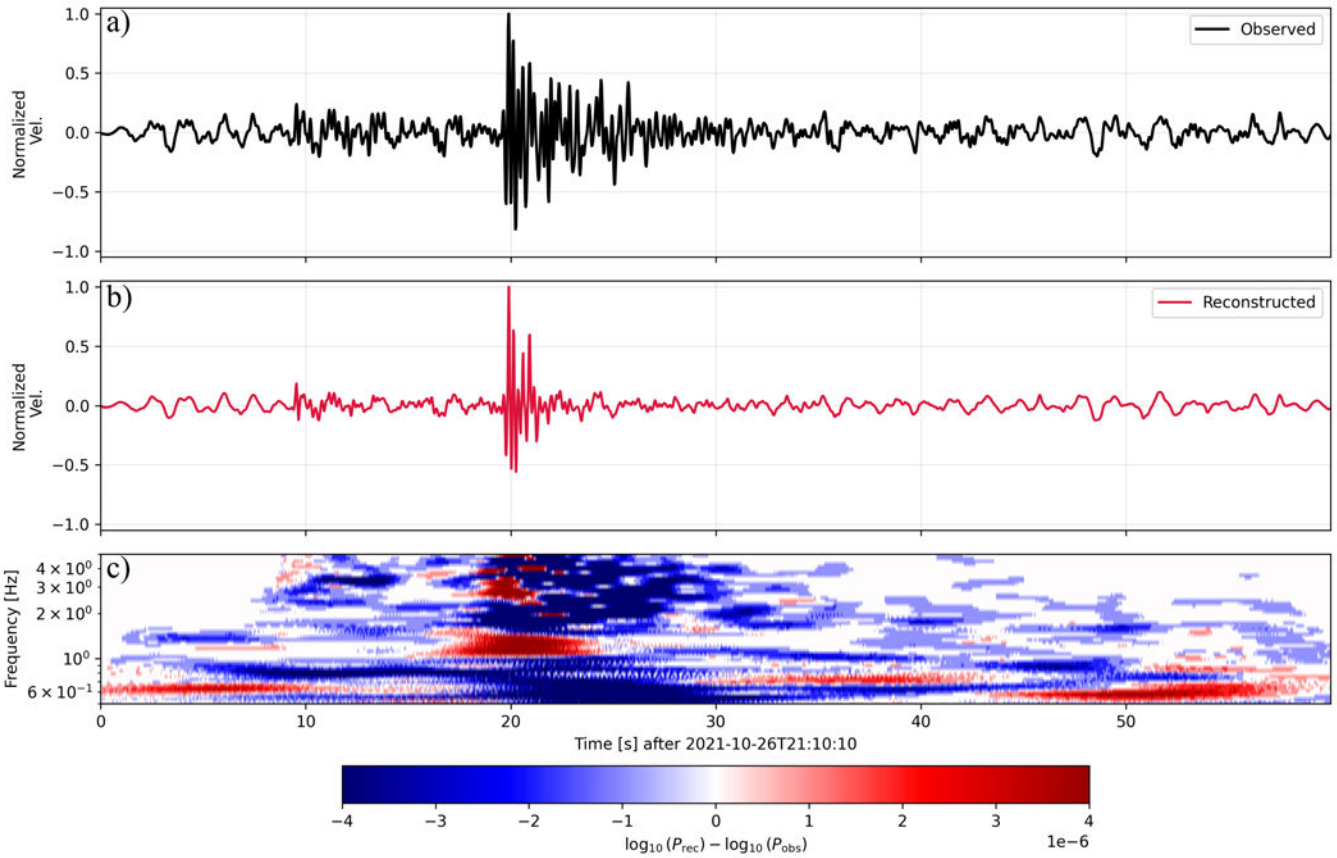


**Figure 6** Time-frequency residuals of the vertical-component at station BPH11 for the candidate event analyzed in the *P*-wave *fk* example. (a) Observed vertical-component waveform. (b) Reconstructed waveform, showing reduced amplitude and morel localized phase structure relative to the true signal. (c) Residual time-frequency representation defined as $\log_{10}(P_{\mathrm{rec}}) - \log_{10}(P_{\mathrm{obs}})$, where red and blue colors indicate amplification and suppression of spectral energy in the reconstruction, respectively.

To illustrate this behavior, we examine reconstruction fidelity in the time-frequency domain. We analyze residual scalograms defined as $\log_{10}(P_{\mathrm{rec}}) - \log_{10}(P_{\mathrm{obs}})$ (Figure 6). Within the *P*-wave arrival window (∼10 seconds), the model amplifies higher-frequency energy above ∼3 Hz while suppressing intermediate frequencies between ∼2 – 5 Hz immediately following the onset. In the *S*-wave window, residuals exhibit a structured redistribution of energy, with amplification near ∼1 Hz, suppression around ∼2 Hz, and renewed amplification between ∼3 – 4 Hz. This is followed by attenuation of higher-frequency content above ∼4 Hz, although amplification of frequencies above 4 Hz is observed immediately following the primary *S*-wave arrival. The model also consistently suppresses *S*-wave coda energy, suggesting a preference for coherent phase arrivals over later scattered energy, which may likely be more complex. At lower frequencies ($< 1$ Hz), energy is attenuated within the main *P*- and *S*-wave windows but amplified at the beginning and end of the time-series, suggesting that the model enhances low-frequency background noise outside of coherent arrivals. These patterns imply that *VASA* emphasizes array coherent phase energy while suppressing less coherent or scattered components of the wavefield. The observed array waveforms used as context can

be seen in Figure S6. Moreover, we show that the reconstructed *P*-wave arrival at BPH11 exhibits higher correlation with neighboring stations across the array than its observed counterpart (Figure S7).

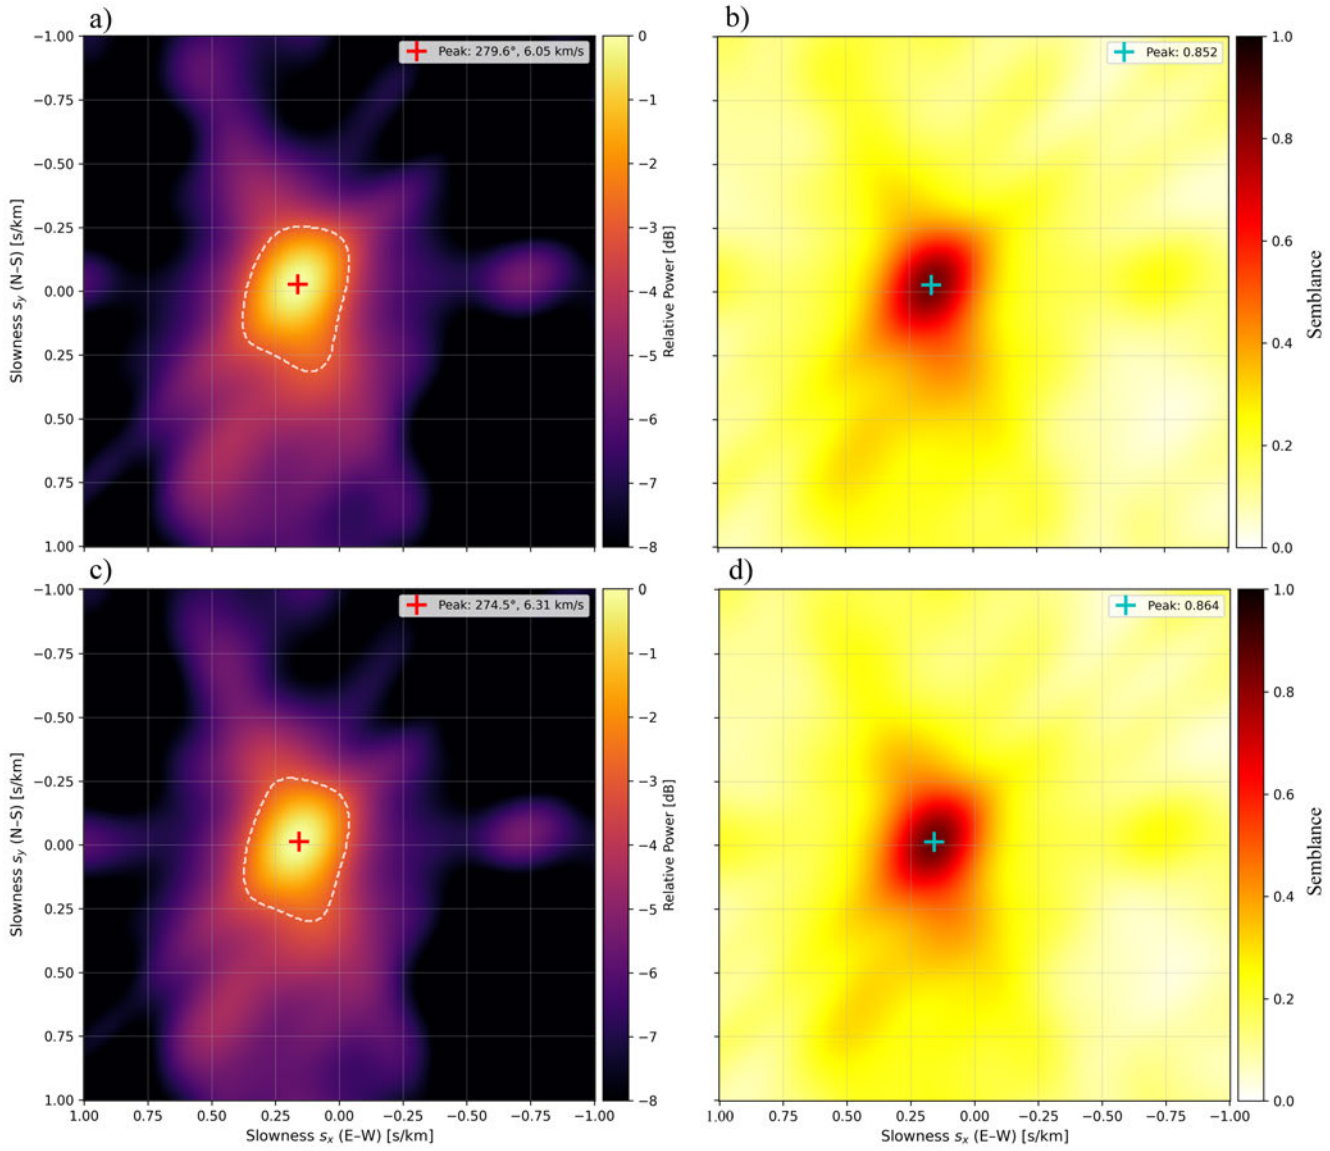


**Figure 7** Analysis of the *P*-wave arrival from Figure 6 and Figure S7 using *fk* processing with a 2-second time window centered on the phase and bandpass filtered between 0.5 – 5 Hz. Ground-truth back azimuth is 274.8°. (a) Relative power for observed waveforms with the peak slowness vector indicated by the red cross and the dashed white line showing the -3 dB contour. (b) Corresponding semblance for the observed data. (c – d) Same as (a – b), but for reconstructed waveforms at the masked sensor (BPH11).

A 2-second window centered on the *P*-wave arrival was used for *fk* analysis to balance phase isolation with sufficient time-frequency support for stable slowness estimation. For the representative BPH11 reconstruction, the reconstructed waveform yields a back azimuth of 274.5°, closely matching the ground-truth value of 274.8° and improving upon the observed estimate of 279.6° (Figure 7). The reconstructed and observed trace velocities are also comparable, indicating that the model preserves the directional and kinematic structure of the *P*-wavefield. Notably, along with the improved directionality estimate, the reconstructed waveform produces a slightly higher semblance value (0.864) than the observed data (0.852), suggesting improved array coherence. This is further supported by the tighter -3 dB contour in the reconstructed *fk* spectrum, indicating a more localized and well-constrained slowness peak. A similar result is observed for the *S*-wave window, where the reconstructed waveform exhibits higher spatial correla-

tion across the array compared to the observed *S*-wave (Figure S8). Additionally, the reconstructed waveform yields a back azimuth of 269.4°, more closely matching the ground-truth value of 264.3° compared to the observed estimate of 273.9° (Figure S9). The reconstructed *fk* spectrum also exhibits a slightly higher semblance, indicating that phase alignment and wavefield structure are preserved for later-arriving phases.

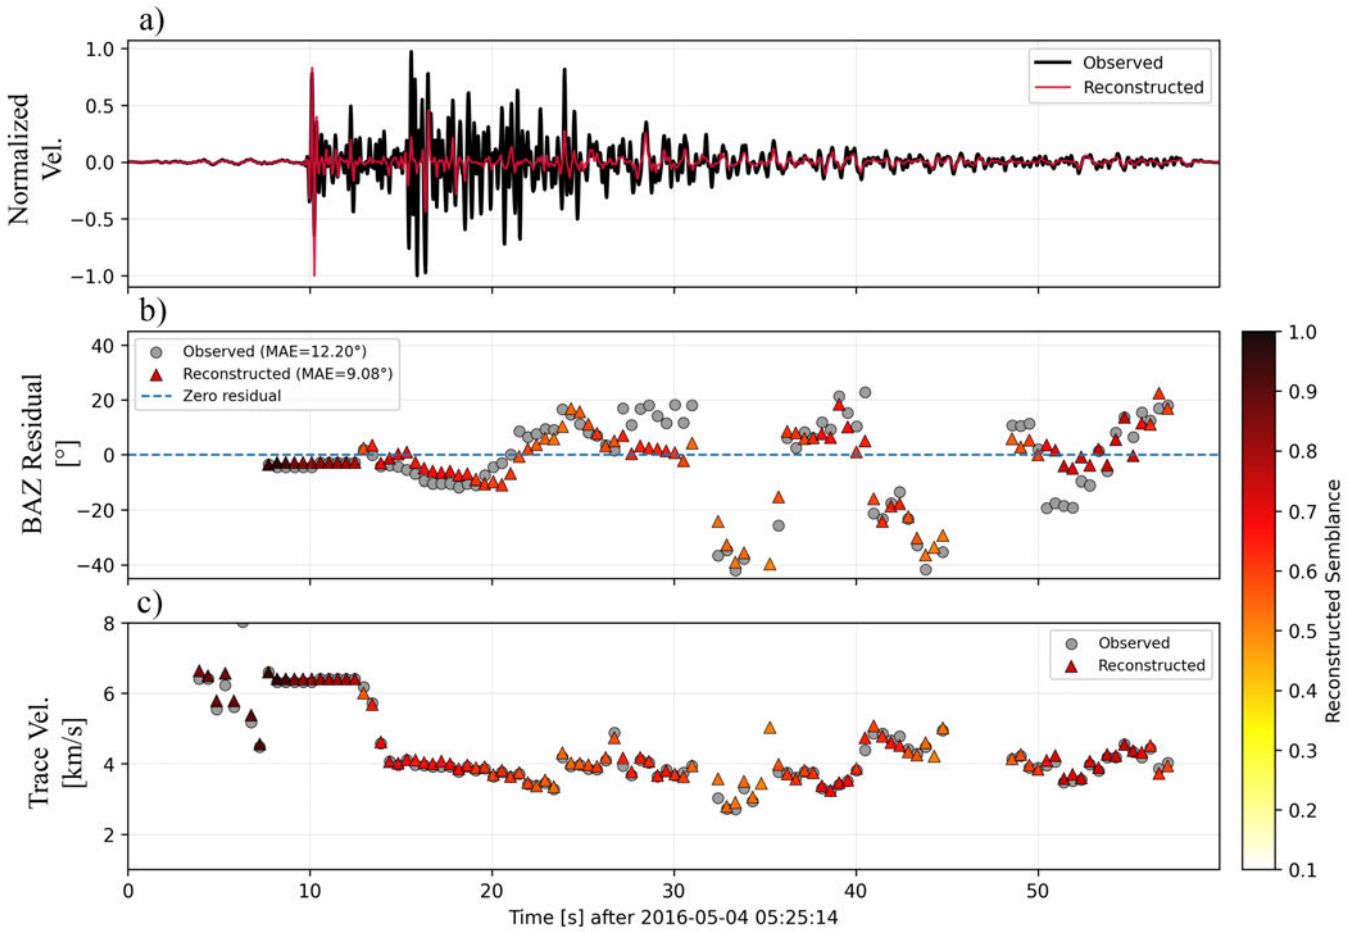


**Figure 8** Sliding-window array processing comparison for a challenging reconstruction case. (a) Observed (black) and reconstructed (red) vertical-component waveform at BPH11, where the reconstruction departs substantially from the true waveform. (b) Back azimuth (BAZ) residuals relative to the ground-truth direction (69.2°), with reconstructed estimates (triangles) showing reduced scatter and lower MAE (9.08°) compared to observed results (12.20°). Colors denote semblance only for the reconstructed array processing results. (c) Trace velocity evolution, demonstrating that despite waveform differences, the reconstruction preserves the dominant propagation characteristics of the wavefield across the array.

To further stress-test *VASA*, we deliberately select an event in which the reconstructed waveform at BPH11 exhibits substantial deviation from the observed signal in the *Z*-component (Figure 8a), with a maximum normalized cross-correlation of only 0.638. This represents a worst-case reconstruction scenario, allowing us to further evaluate whether physically meaningful wavefield properties, such as directionality and trace velocity, are preserved even when waveform similarity is poor. We extend this analysis beyond isolated phase windows and apply sliding-window *fk* analysis across the full duration of the event, enabling a time-resolved assessment of how well the model preserves wavefield structure and propagation characteristics across all phases. Directionality estimates are plotted as back azimuth residuals in Figure 8b. For each sliding *fk* window, we estimate the back azimuth from the observed and reconstructed array data, subtract the ground-truth back azimuth, and wrap the difference onto the circular interval $[-180°, 180°)$. Therefore, values near zero indicate that the estimated propagation direction agrees with the known

source direction, while positive or negative values indicate angular deviations from the expected back azimuth. Associated trace velocities for the observed and reconstructed waveforms are plotted in Figure 8c. Lastly, we report the mean absolute error (MAE) separately for the observed and reconstructed results. The MAE is computed as the average absolute back azimuth residual across valid sliding windows, where valid windows are defined independently for each dataset as those with semblance $> 0.5$ and $|\mathrm{BAZ}_{\mathrm{residual}}| \leq 45^{\circ}$.

Sliding-window *fk* results further demonstrate that the reconstructed wavefield not only preserves, but in many cases improves, directional consistency relative to the observed data. This is reflected in the reduction of back azimuth MAE from $12.20^{\circ}$ for the observed waveforms to $9.08^{\circ}$ for the reconstruction (Figure 8b). Notably, the reconstructed estimates exhibit improved agreement with the ground-truth back azimuth across multiple segments of the event. This includes the pre-*P*-wave interval, the *S*-wave window (∼15 – 19 s), and the post-*S*-wave coda where coherent arrivals persist (∼26 – 40 s and beyond). In these intervals, the reconstructed back-azimuth residuals are consistently closer to zero and exhibit reduced scatter compared to the observed estimates. While localized intervals near ∼20 s and ∼25 s show slightly larger deviations in the reconstruction, the overall behavior indicates that *VASA* enhances the stability and coherence of directional estimates across the entirety of the event. In the trace velocity estimates (Figure 8c), reconstructed values closely follow the observed propagation speeds, with slight overestimation during some *P*- and *S*-wave segments and minor underestimation in the coda. Overall, these results show that *VASA* retains the fundamental propagation physics of the seismic wavefield, including phase alignment, directionality, and kinematics, even in difficult reconstruction scenarios.

## 4.3 Interpretability Analysis

To investigate how *VASA* distributes information across the array during reconstruction, we analyze attention weights from its Transformer blocks. For each validation event, every sensor is masked once, and attention is extracted for the masked sensor tokens. These token-level attention values are then collapsed across the three components to produce a sensor-level attention matrix, where rows represent the masked query sensor and columns represent contributing key sensors. The matrix is row-normalized, such that each row sums to one, allowing the relative contribution of each sensor to the reconstruction of a given masked sensor to be interpreted directly. The resulting attention matrix shows that the model does not rely uniformly on all sensors (Figure 9). Instead, BPH09 and BPH11 receive consistently elevated attention across many query sensors, suggesting that these more spatially isolated stations provide non-redundant constraints on the array wavefield. This pattern is notable because BPH11 is also the most difficult sensor to reconstruct, indicating that reconstruction difficulty and informational value are not equivalent. In other words, isolated sensors may be harder to infer from the rest of the array, while simultaneously contributing disproportionately useful information for reconstructing other stations.

Since attention weights alone do not conclusively establish sensor importance, we further test this interpretation using sensor-removal experiments. In these experiments, reconstruction of a target station is evaluated under two conditions: a baseline case in which only the target station is masked, and a comparison case in which the target station is masked and an additional context station is zeroed out, effectively removing it as a contributing sensor. We focus on reconstruction of BPH09 and assess the effect of removing either BPH11, which received consistently high attention in the interpretability analysis, or BPH02, a less isolated station that received comparatively low attention.

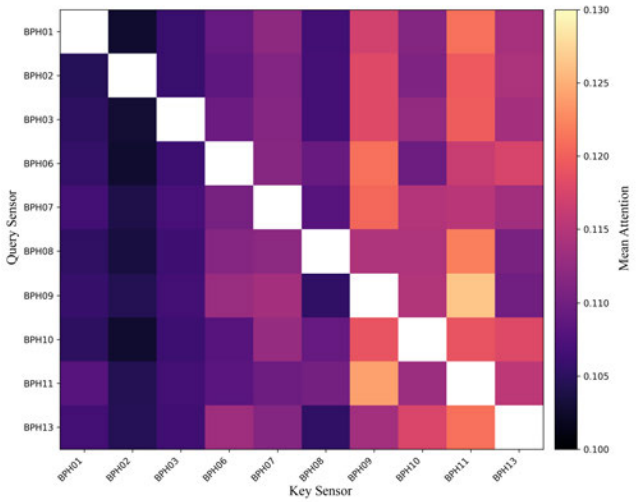


**Figure 9** Sensor-level attention weights for masked-sensor reconstruction on the validation set. Each row corresponds to a masked (query) sensor, and each column indicates the relative contribution of a context (key) sensor to its reconstruction. Attention values are row-normalized, such that each row sums to one and can be interpreted as a distribution of importance across sensors. Warmer colors indicate greater reliance on a given sensor.

Reconstruction degradation is quantified relative to the baseline case using maximum normalized cross-correlation, RMSE, best lag, RMS ratio, and peak-amplitude ratio.

Removing BPH11 substantially degrades reconstruction of BPH09, reducing the mean maximum cross-correlation from 0.869 to 0.763 and increasing RMSE from $8.17 \times 10^{-8}$ to $1.00 \times 10^{-7}$. The amplitude metrics also decrease, with the RMS ratio dropping from 0.889 to 0.636 and the peak ratio from 1.133 to 0.727. This degradation occurs consistently across components, with cross-correlation worsening in more than 99% of event-component examples. Removing BPH02 also degrades reconstruction, reducing mean cross-correlation to 0.778 and increasing RMSE to $1.02 \times 10^{-7}$. However, the drop in cross-correlation is larger when BPH11 is removed than when BPH02 is removed, indicating that BPH11 provides especially important contextual information for preserving waveform similarity in reconstruction of BPH09. At the same time, removal of either station produces substantial degradation in amplitude-based metrics, suggesting that both sensors contribute useful context, although the stronger correlation loss associated with BPH11 is consistent with the attention-based interpretation that the most spatially isolated stations provide uniquely informative, non-redundant constraints on the array wavefield.

# 5 Discussion

We leverage *VASA* to learn representations of the seismic wavefield using local earthquakes recorded by a regional array in southern California (PFO). The model achieves near perfect reconstruction performance across waveform similarity and amplitude-based metrics at well-connected sensors (e.g., BPH01, BPH02, BPH03), indicating that the learned representation captures the dominant spatiotemporal structure of the wavefield in regions of dense observational support. Attention-based analysis also shows that the model relies comparatively more on spatially isolated sensors for the general task of three-component waveform reconstruction. However, the central question is whether *VASA* learns more than a sophisticated form of data interpolation. To address this, we focus our analysis on the most difficult-to-reconstruct sensor, where the reconstructed waveform demonstrably deviates from the

observed waveform. This provides a stringent test of what information the model prioritizes when reconstruction is relatively poor. Despite reduced waveform similarity at this sensor, array processing results demonstrate that the reconstructed signals preserve physically meaningful characteristics of the wavefield, including arrival timing, directionality, trace velocity, and spatial coherency across the array. Notably, Klinge et al. (2026) also reported improvements in beamforming-derived directionality using a Virtual Seismic Array, but that study was conducted in a fundamentally different setting, where the database consisted of ambient noise wavefields in the narrow secondary microseism band (0.1 – 0.25 Hz). Overall, this behavior indicates that *VASA* learns an underlying representation of wavefield structure governed by propagation physics. In particular, the model appears to prioritize spatially coherent features associated with propagation phases, while deemphasizing fine-scale waveform variability that does not significantly influence array-level observables. These results suggest that the learned latent space encodes a physically meaningful description of the seismic wavefield across regional array scales, enabling robust recovery of propagation characteristics even in challenging reconstruction scenarios.

A key downstream objective of *VASA* is not only accurate reconstruction within a single array, but robust generalization to new regions, array geometries, and event types. Several design choices support this goal. First, the Transformer incorporates a distance-informed spatial bias, conditioning cross-sensor attention on inter-station geometry. Unlike prior approaches that learn station relationships purely from data, this explicitly encodes a physics-consistent notion of spatial coherency, theoretically allowing the model to adapt to arrays with different layouts without relying on station-specific correlations. Second, the use of a global robust scaling scheme based on a MAD normalization combined with an *asinh* transform enables the model to accommodate large dynamic range in amplitudes while remaining insensitive to outliers and region-specific magnitude distributions. This encourages the model to learn waveform structure rather than dataset-specific magnitude distributions. Finally, magnitude-balanced training through oversampling mitigates the common tendency of reconstruction models to underestimate large-amplitude signals by ensuring sufficient representation of higher-magnitude events during training. These elements promote a representation that is governed by propagation physics and waveform structure. While the current evaluation is conducted in an in-domain setting, where training and testing events occupy similar source regions (Figure S1), this provides a controlled environment to assess whether the model learns physically meaningful wavefield structure. We expect these methodological choices to enhance the model's transferability across different seismic settings and event types, a hypothesis we elaborate on further in the following section.

To further assess the robustness of the learned representation, we examine a series of challenging scenarios designed to probe whether *VASA* preserves coherent wavefield structure, while disregarding features that are not physically meaningful at the array level. First, in the case of multiple overlapping events within a single window, the model successfully reconstructs both arrivals, indicating that it captures the superposition of propagating wavefields rather than assuming a single dominant phase (Figure S10). This behavior suggests that the model encodes the underlying linear structure of seismic wave propagation. Second, for examples in which the primary arrival is temporally offset from the nominal reference position, the model maintains high-fidelity reconstruction despite the shift, demonstrating that it is not positionally biased or tied to fixed temporal priors but instead learns relative spatiotemporal structure such as moveout, phase alignment, and inter-sensor coherency (Figure S11). Finally, in the presence of im-

pulsive, component-specific noise transients, the model does not reproduce these features and instead preserves the coherent signal across the array and sensor, even at the expense of reduced waveform similarity metrics (Figure 10). This selective suppression of incoherent energy further indicates that the model prioritizes physically meaningful wavefield characteristics over localized, non-wavefield artifacts. These results provide strong evidence that *VASA* learns an implicit representation of the seismic wavefield governed by propagation physics, which likely underpins its robustness and supports its potential for transferability to new geologic regions, arrays, and event types.

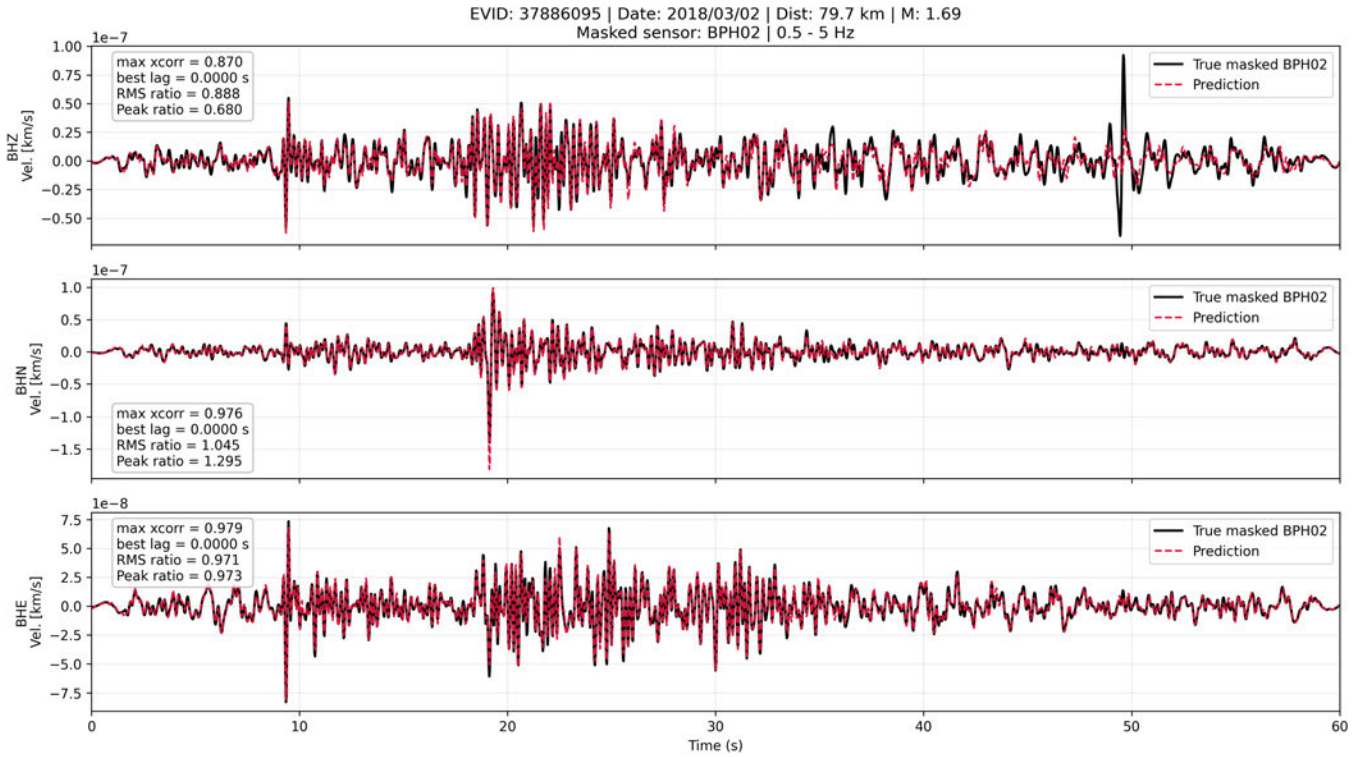


**Figure 10** Same as Figure 5, but for an event window containing an impulsive noise transient isolated to the *Z*-component at ∼50 seconds. The reconstruction does not reproduce this noise feature, instead preserving the coherent wavefield signal despite a resulting decrease in waveform correlation.

To assess the frequency dependence of reconstruction performance, we trained an identical model on a broader 0.5 – 10 Hz dataset. While the model exhibits stable optimization behavior and converges smoothly (Figure S12), the training history shows a persistent gap between training and validation loss, with training loss remaining higher. This behavior likely reflects the increased difficulty of the training objective due to stochastic masking and augmentation, rather than classical overfitting. This suggests that the current training pipeline may not be optimally configured for reconstructing higher-frequency (5 – 10 Hz) waveform content. Additionally, overall validation reconstruction quality is reduced relative to the 0.5 – 5 Hz case, with lower mean cross-correlation values ($\sim$0.69 – 0.76 across components) and increased variability in amplitude metrics (Figure S12b-d). Notably, reconstruction performance is not uniform across components. The *Z*-component is recovered more accurately than the horizontal components, which exhibit systematically lower correlation and greater amplitude variability (Figure S13). This disparity likely reflects differences in wavefield coherency and signal-to-noise characteristics, with vertical motion dominated by more coherent body wave energy, whereas horizontal components contain more complex scattering, site effects, and polarization variability at higher frequencies.

This trend is further reflected at the sensor level (Figure S14a), where the maximum normalized cross-correlation degrades most strongly at more isolated stations ($\sim$0.55 – 0.65), while staying much higher at the tightly clustered sensors ($\sim$0.90). In addition, the isolated sensors exhibit pronounced amplitude deviations from 1.0, including systematic RMS ratio underestimation (Figure S14b). Furthermore, the peak ratio distributions indicate a tendency to overestimate high-frequency amplitudes more broadly, but most dramatically observed at the BPH08 sensor (Figure S14c). These 0.5 – 10 Hz training and reconstruction quality results underscore an important limitation of the present study. Although the 0.5 – 5 Hz band captures the dominant coherent wavefield structure, significant signal energy exists above 5 Hz, particularly at local to regional distances. This higher frequency content is typically leveraged for critical monitoring tasks such as event discrimination between anthropogenic and natural sources. The combination of reduced component-dependent performance, waveform-based similarity metrics, and spatially variable amplitude bias at higher frequencies suggests that additional model capacity, revised inductive biases, or alternative training strategies will be required to fully capture the broadband wavefield behavior in this regime.

# 6 Conclusion

In this study, we introduce *VASA*, a U-Net-inspired convolutional-Transformer architecture designed to implicitly learn seismic wavefield structure across a regional array through self-supervised reconstruction. Rather than framing the problem as interpolation, we demonstrated that the model captures physically meaningful, array-consistent representations of the wavefield. This is evidenced by array processing validation, where reconstruction yields slowness and back azimuth estimates that more closely match ground-truth in challenging reconstruction scenarios, indicating preservation of phase alignment and propagation characteristics. Additional analyses further show that the model generalizes across complex waveform conditions. In particular, it successfully reconstructs multi-event windows, remains robust to temporal offsets in arrival structure, and selectively suppresses component-specific transient noise that is not coherent across the array. We attribute the latter to the component-specific multi-branch design of the neural network. Attention-based interpretability reveals that more isolated sensors, whose observations im-

pose stronger geometric constraints on the wavefield, contribute disproportionately more to reconstruction, while more tightly clustered sensors contribute less, and instead play a comparatively more redundant role. This interpretability analysis quantifies how informational value is distributed across sparse regional-scale arrays, with direct implications to the future design and deployment of these types of sensor configurations. All in all, these results suggest that *VASA* has learned an implicit representation of the seismic wavefield and that it prioritizes physically consistent array-level features such as arrival timing, wavefront coherency, and propagation geometry over simply reproducing waveform amplitudes.

Looking forward, we aim to extend this framework toward a robust deployment and generalization across diverse seismic settings. The primary goal of the present study was to establish that the proposed neural network architecture can learn physically meaningful seismic wavefield structure within a well-constrained regional setting. Therefore, we defer explicit generalization experiments to future work, as both new event types and new geologic regions introduce fundamentally different wavefield characteristics. In particular, differing source mechanisms generate distinct radiation patterns and polarization behavior, while variations in regional subsurface structure and array geometry alter propagation effects and spatial coherency. Furthermore, a key next step is the incorporation of a multi-sensor masking strategy during training, allowing the model to explicitly learn reconstruction under arbitrary sensor-loss scenarios. These capabilities are already supported by *VASA* but not fully exploited in the present study. As mentioned previously, we will also investigate transferability to new event types (e.g., explosions) and geologic regions, where fine-tuning is likely required. Importantly, several methodological choices introduced here, including distance-informed attention, robust amplitude scaling, and training strategies that promote learning of underlying wavefield structure, are expected to facilitate this generalization. Future work will therefore focus on systematically evaluating these capabilities in a deployment-oriented setting, with the goal of establishing *VASA* as a broadly applicable framework for physics-informed seismic wavefield modeling.

## Data and code availability

The Cardinal software package can be found at Arrowsmith (2020). Tensorflow was used as the deep learning framework for this study Abadi et al. (2016) and Obspy was used as the framework for geophysical analysis Beyreuther et al. (2010). Associated seismic data, along with link to supplemental code, can be found at Ronac Giannone (2026).

## Competing interests

The authors declare no conflicts of interest relevant to this study.

## References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: A system for large-scale machine learning. *12th USENIX symposium on operating systems design and implementation*, pages 265–283, 2016. doi: 10.48550/arXiv.1605.08695.

Abedi, M. M. and Pardo, D. A Multidirectional Deep Neural Network for Self-Supervised Reconstruction of Seismic Data. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 2022. doi: 10.1109/TGRS.2022.3227212.

Abedi, M. M., Pardo, D., and Alkhalifah, T. A. Ensemble Deep Learning for Enhanced Seismic Data Reconstruction. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–11, 2024. doi: 10.13039/501100011033.

Aki, K. and Richards, P. G. *Quantitative Seismology*. University Science Books, Sausalito, California, 2nd edition, 2002.

Arrowsmith, S. J. Cardinal [Software]. https://github.com/sjarrowsmith/cardinal, 2020.

Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., and Wassermann, J. ObsPy: A Python Toolbox for Seismology. *Seismological Research Letters*, 81:530–533, 2010. doi: 10.1785/gssrl.81.3.530.

Chen, Y., Bai, M., and Chen, Y. Obtaining free USArray data by multi-dimensional seismic reconstruction. *Nature Communications*, 10, 2019. doi: 10.1038/s41467-019-12405-0.

Christie, D. R. and Campus, P. *The IMS infrasound network: Design and establishment of infrasound stations*, pages 29–75. Springer Netherlands, 2010. doi: 10.1007/978-1-4020-9508-5_2.

Darrh, A., Poppeliers, C., and Preston, L. Azimuthally Dependent Seismic-Wave Coherence at the Source Physics Experiment Large-N Array. *Bulletin of the Seismological Society of America*, 109(5):1935–1947, 2019. doi: 10.1785/0120180296.

Gaharwar, A., Kulkarni, P. P., Dickey, J., and Shah, M. Xi-Net: Transformer based Seismic Waveform Reconstructor. In *Proceedings - International Conference on Image Processing, ICIP*, pages 2725–2729. IEEE Computer Society, 2023. doi: 10.1109/ICIP49359.2023.10222465.

Gibbons, S. J., Kværna, T., and Ringdal, F. Considerations in Phase Estimation and Event Location Using Small-aperture Regional Seismic Arrays. *Pure and Applied Geophysics*, 167(4):381–399, 2010. doi: 10.1007/s00024-009-0024-1.

He, T., Wu, B., and Zhu, X. Seismic Data Consecutively Missing Trace Interpolation Based on Multistage Neural Network Training Process. *IEEE Geoscience and Remote Sensing Letters*, 19, 2022. doi: 10.1109/LGRS.2021.3089585.

Huang, H., Wang, T., Cheng, J., Xiong, Y., Wang, C., and Geng, J. Self-Supervised Deep Learning to Reconstruct Seismic Data with Consecutively Missing Traces. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 2022. doi: 10.1109/TGRS.2022.3148994.

Klinge, J., Schippkus, S., Walda, J., Hadziioannou, C., and Gajewski, D. Predictive modelling of seismic wave fields: learning the transfer function using encoder–decoder networks. *Geophysical Journal International*, 240:1611–1621, 2025. doi: 10.1093/gji/ggaf004.

Klinge, J., Schippkus, S., and Hadziioannou, C. Unlocking the potential of single stations to replace seismic arrays. *Geophysical Journal International*, 245, 2026. doi: 10.1093/gji/ggag070.

Koper, K. D. The importance of regional seismic networks in monitoring Nuclear Test-Ban Treaties. *Seismological Research Letters*, 91(2A): 573–580, 2020. doi: 10.1785/0220190160.

Li, L. and Schwarz, B. Coherent Wavefield Reconstruction Improves Event Location With Dense Seismic Arrays. *IEEE Geoscience and Remote Sensing Letters*, 21:2024, 2024. doi: 10.5281/zenodo.13120750.

Meng, F., Fan, Q., and Li, Y. Self-Supervised Learning for Seismic Data Reconstruction and Denoising. *IEEE Geoscience and Remote Sensing Letters*, 19, 2022. doi: 10.1109/LGRS.2021.3068132.

Muir, J. B. and Zhan, Z. Seismic wavefield reconstruction using a pre-conditioned wavelet-curvelet compressive sensing approach. *Geophysical Journal International*, 227:303–315, 2021. doi: 10.1093/gji/ggab222.

Ronac Giannone, M. Virtual Array Synthesis Architecture, May 2026. doi: 10.5281/zenodo.20219324.

Ronac Giannone, M., Arrowsmith, S. J., and Silber, E. A. Cardinal: Seismic and Geoacoustic Array Processing. *Seismological Research Letters*, 97(1):487–500, 2025. doi: 10.1785/0220240434.

649 Ronac Giannone, M., Arrowsmith, S., and Averbuch, G. Unsupervised Deep Representation Learning for Infrasound Phase Identification.
650 *Authorea*, February 2026. doi: 10.22541/au.177032791.13330612/v2.

651 Ronneberger, O., Fischer, P., and Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Navab, N., Hornegger,
652 J., Wells, W. M., and Frangi, A. F., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241.
653 Springer International Publishing, Cham, 2015. doi: 10.1007/978-3-319-24574-4_28.

654 Rost, S. and Thomas, C. Array seismology: Methods and applications. *Reviews of Geophysics*, 40, 2002. doi: 10.1029/2000RG000100.

655 Rost, S. and Thomas, C. Improving seismic resolution through array processing techniques. *Surveys in Geophysics*, 30:271–299, 2009.
656 doi: 10.1007/s10712-009-9070-6.

657 Shi, P., Seydoux, L., and Poli, P. Unsupervised Learning of Seismic Wavefield Features: Clustering Continuous Array Seismic Data During
658 the 2009 L'Aquila Earthquake. *Journal of Geophysical Research: Solid Earth*, 126(1):e2020JB020506, 2021. doi: 10.1029/2020JB020506.

659 Stump, B. W., Hayward, C., Golden, P., Park, J., Kubacki, R., Cain, C., Arrowsmith, S., Taylor, M. H. M., Jeong, S. J., Ivey, T., MacPhail, M., Pace,
660 C. C., Jeon, J. S., Che, I. Y., Kim, K., Kim, B. I., Kim, T. S., Shin, I. C., and Jun, M. S. Seismic and Infrasound Data Recorded at Regional
661 Seismoacoustic Research Arrays in South Korea from the Six DPRK Underground Nuclear Explosions. *Seismological Research Letters*, 93:
662 2389–2400, 2022. doi: 10.1785/0220220009.

663 Wang, B., Zhang, N., Lu, W., and Wang, J. Deep-learning-based seismic data interpolation: A preliminary result. *Geophysics*, 84:V11–V20,
664 2019. doi: 10.1190/geo2017-0495.1.

665 Yu, J. and Wu, B. Attention and Hybrid Loss Guided Deep Learning for Consecutively Missing Seismic Data Reconstruction. *IEEE Transactions*
666 *on Geoscience and Remote Sensing*, 60, 2022. doi: 10.1109/TGRS.2021.3068279.

**Supplement to:**

# Learning Seismic Wavefield Structure from Regional Arrays with Self-Supervised Deep Learning

Description:

Figures outline train/test event location distributions, neural network training results, per component and per sensor waveform-based reconstruction metrics, and array data used as context for reconstruction. Additionally, we provide spatial cross-correlation results between observed and reconstructed sensors, and *fk* results between observed and reconstructed waveforms. We also demonstrate the model's capability at reconstructing surface waves, as well as unique event windows which contain multiple overlapping events and offsets in the nominal reference time, where the initial *P*-wave deviates substantially from its predicted arrival time. Moreover, we demonstrate the model's ability at ignoring impulsive noise segments during reconstruction. Finally, we report on the training results and reconstruction quality metrics, across both components and sensors, for the 0.5 – 10 Hz train/test dataset.

Figures:

Figures S1 – S14

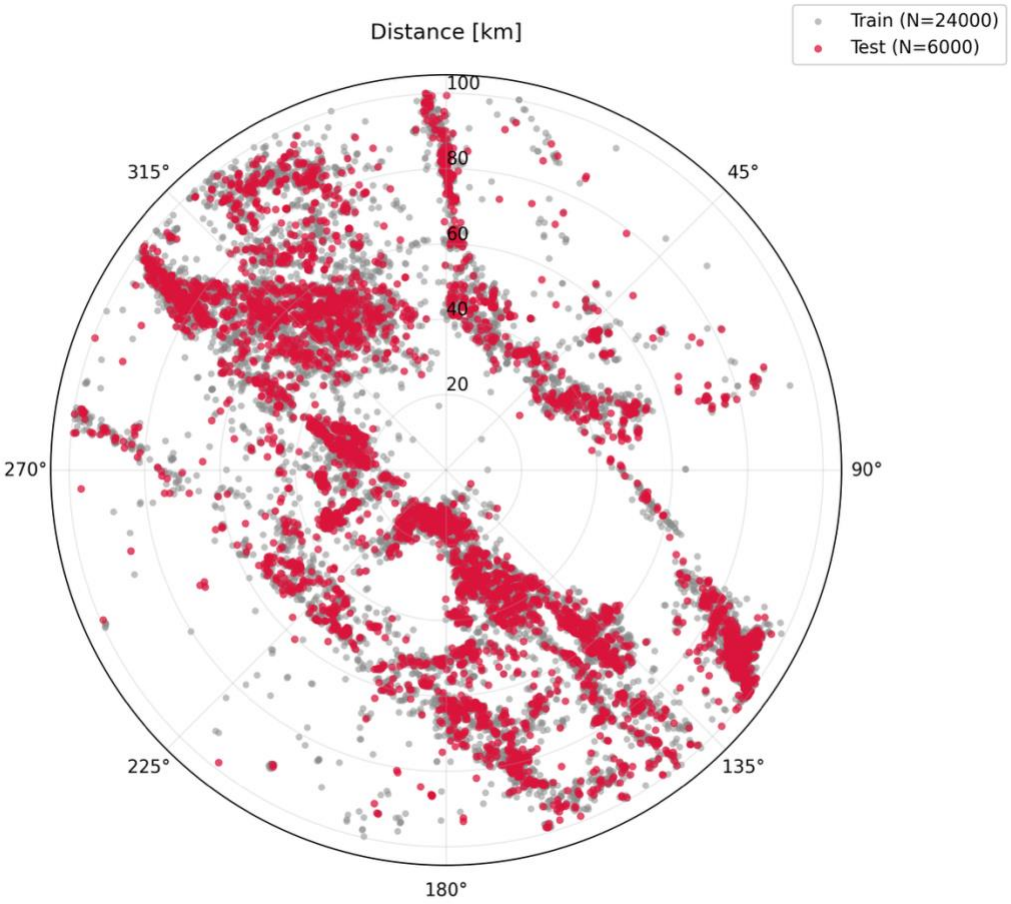


Figure S1: Spatial distribution of training and testing events used in this study. Training events (grey) and testing events (red) are shown as a function of back azimuth and source-receiver distance. While individual events are unique to each partition, the two sets occupy overlapping source regions.

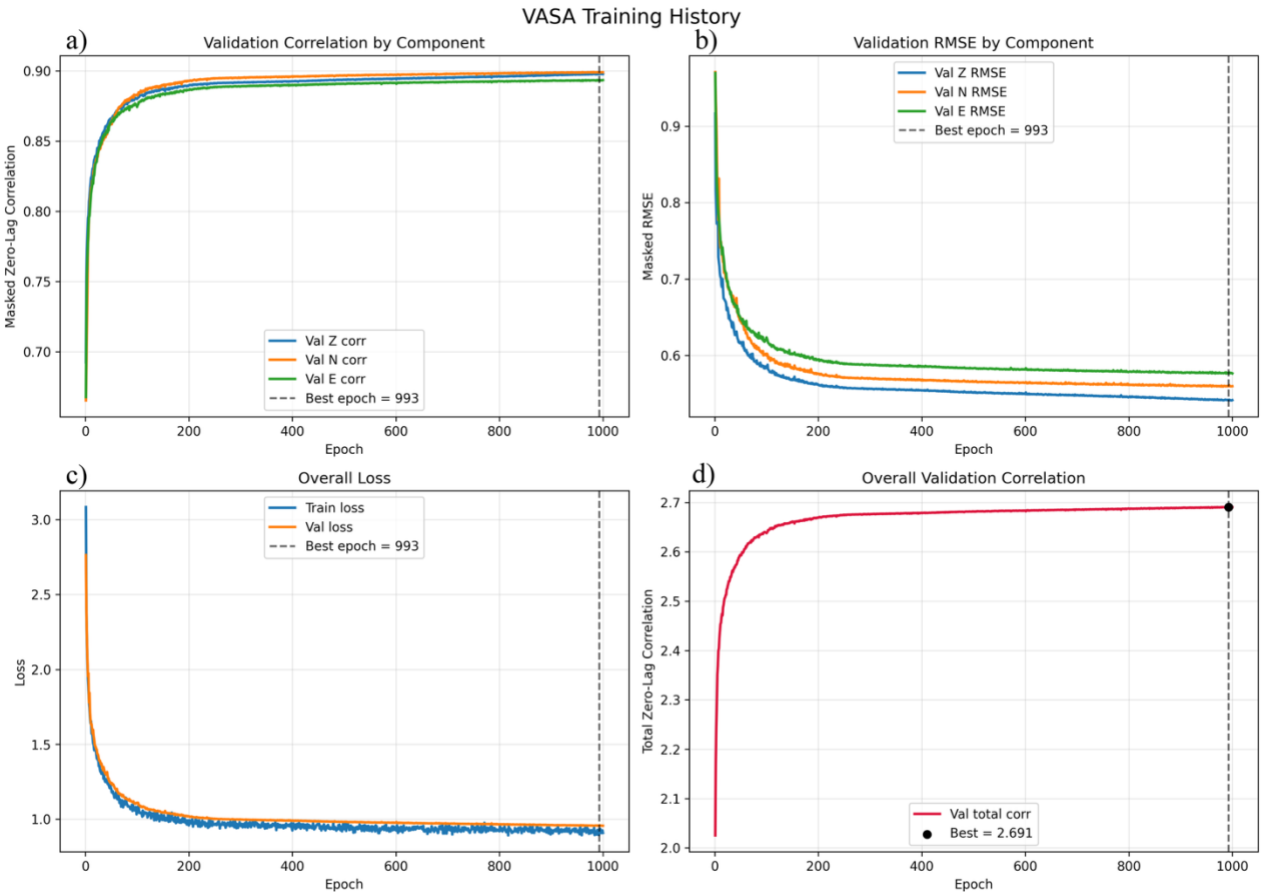


Figure S2: Training history of the *VASA* model under masked single-sensor reconstruction. (a) Validation zero-lag correlation by component, showing rapid convergence and stable performance near 0.90 for all components. (b) Corresponding masked RMSE, which decreases consistently throughout training, with the *Z*-component exhibiting the lowest error. (c) Overall training and validation loss, demonstrating stable optimization with no strong evidence of overfitting. (d) Total validation correlation (summed across components), used for model selection, with the optimal model identified at epoch 993 (black marker) achieving a value of 2.691. Dashed vertical lines indicate the selected epoch.

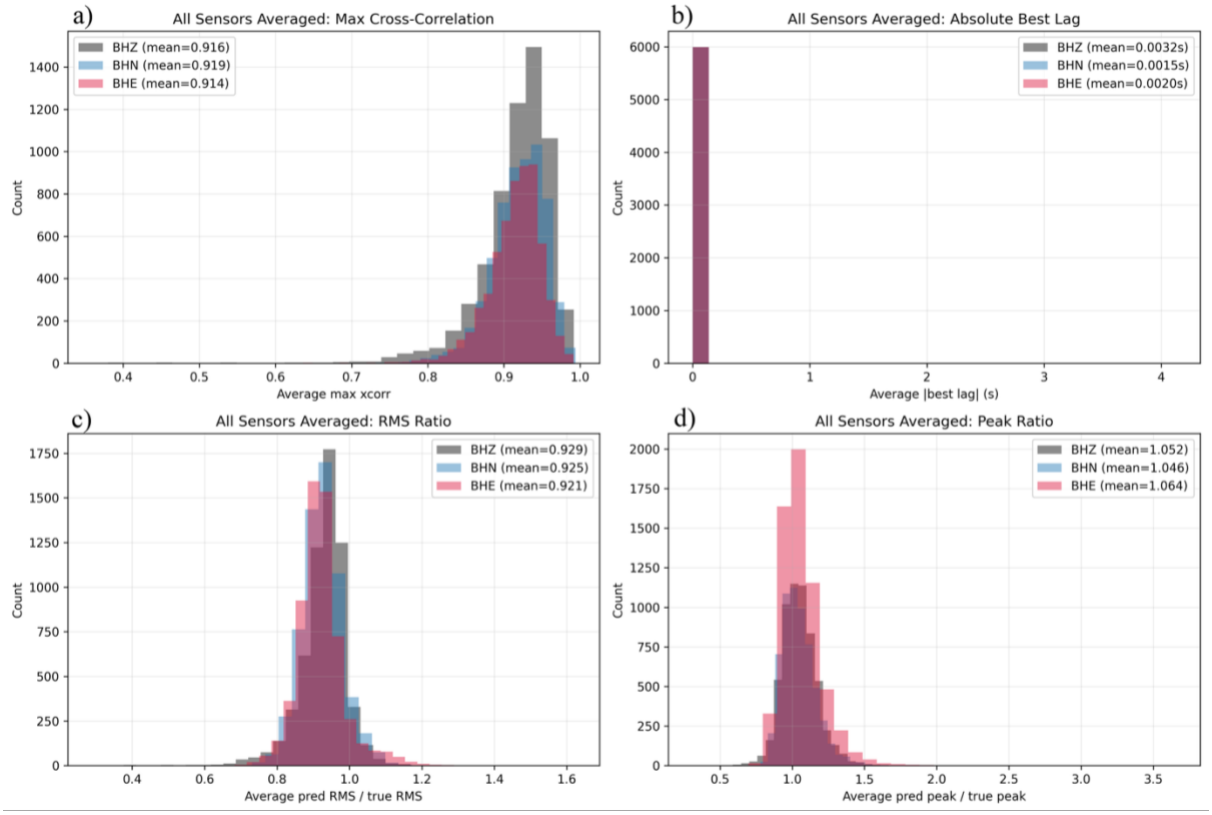


Figure S3: Distributions of reconstruction performance metrics averaged across all sensors. (a) Maximum normalized cross-correlation shows consistently high waveform similarity across components. (b) Best-lag values are sharply centered near zero, indicating accurate preservation of arrival timing. (c) RMS ratios suggest that overall waveform energy is well recovered, with a slight underestimation across components. (d) Peak amplitude ratios are close to unity, demonstrating accurate recovery of dominant waveform amplitudes. These results indicate that there is little to no component bias during waveform reconstruction.

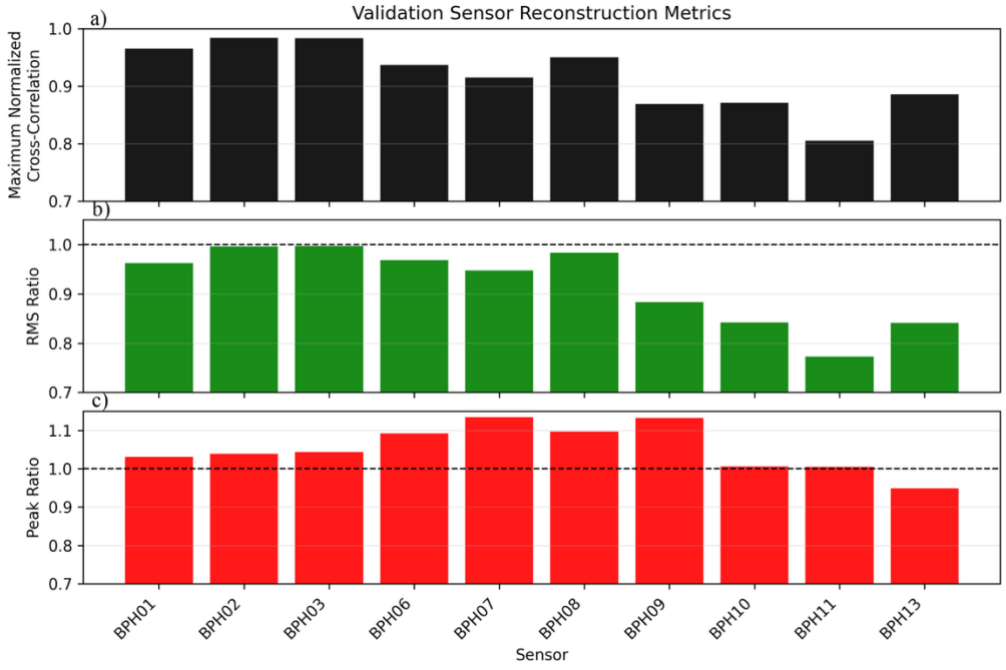


Figure S4: Reconstruction performance by sensor, averaged across all events and components. (a) Maximum normalized cross-correlation, (b) RMS ratio (predicted/true), and (c) peak amplitude ratio for each station.

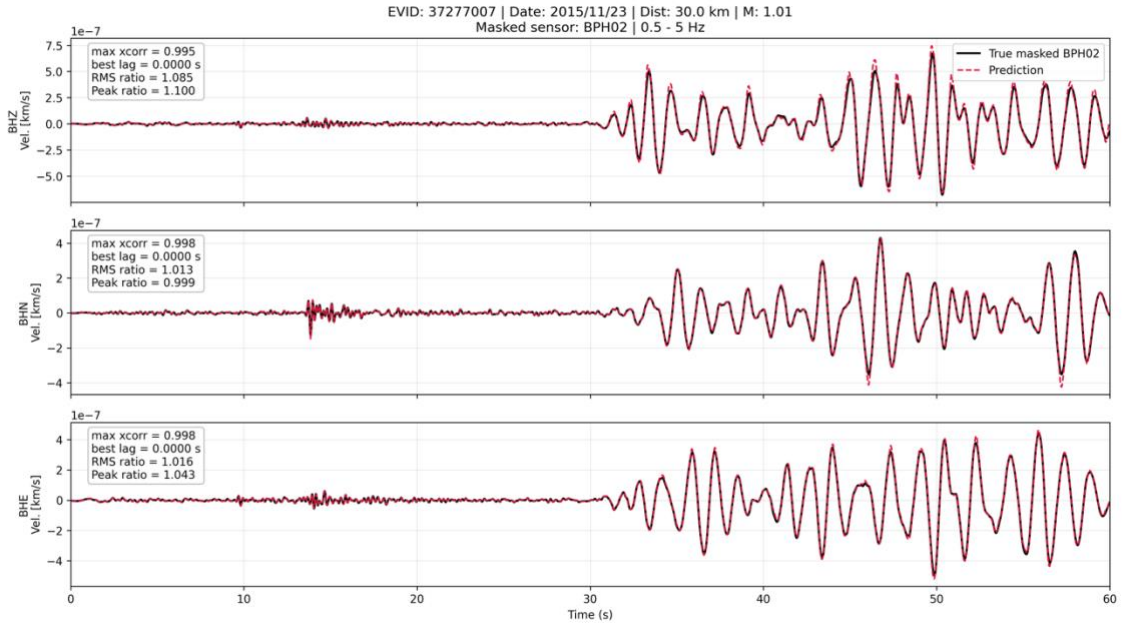


Figure S5: Observed and reconstructed three-component waveforms for a representative event with clear surface wave structure at BPH02. Black lines show the true waveform at the masked sensor, while red dashed lines indicate the model reconstruction. Reconstruction quality metrics follow the same description of Figure 5.

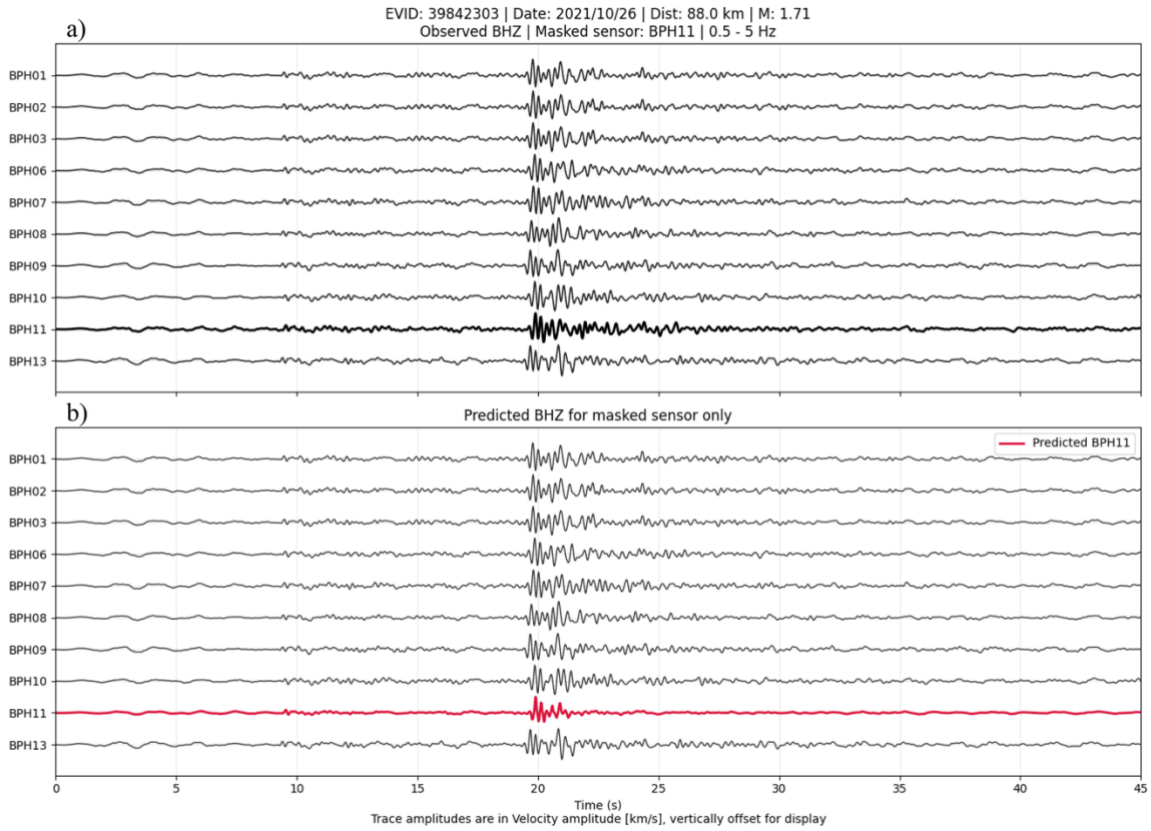


Figure S6: Observed and reconstructed array waveforms for the event used in *P*-wave *fk* analysis. (a) Observed vertical-component data with BPH11 (bold trace) designated as the masked sensor. (b) Reconstruction of BPH11 (red) using the remaining sensors as context. Despite degraded waveform fidelity at BPH11, the reconstructed signal preserves the primary arrival structure, enabling physically meaningful array processing in subsequent analyses.

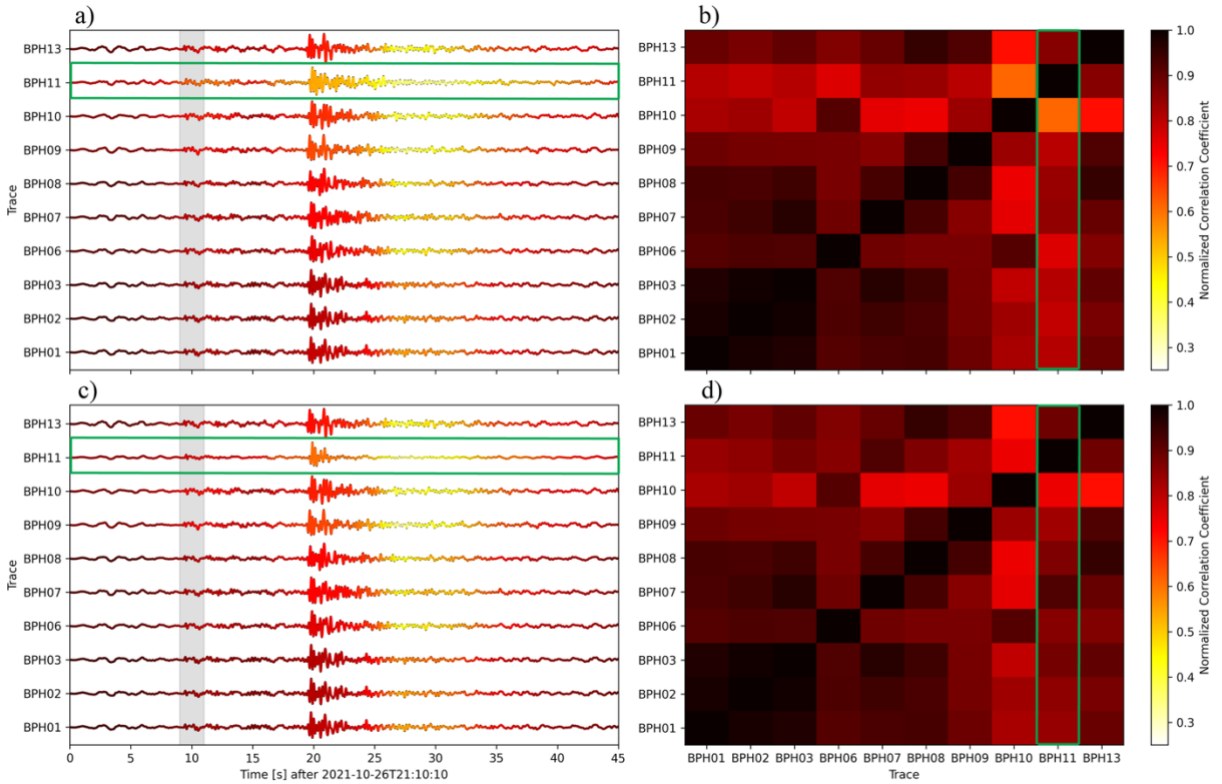


Figure S7: Array waveform structure and inter-station similarity for the *P*-wave example. (a) Observed vertical-component waveforms across the array, with the shaded region indicating the *P*-wave window used for *fk* analysis and cross-correlation. (b) Corresponding inter-station normalized cross-correlation matrix computed over the *P*-wave window. (c–d) Same as (a–b), but for the reconstructed waveform at the masked sensor (BPH11). Green boxes highlight the reconstructed sensor. Reconstruction enhances inter-station correlation structure, indicating that the spatial coherence of the wavefield is maintained.

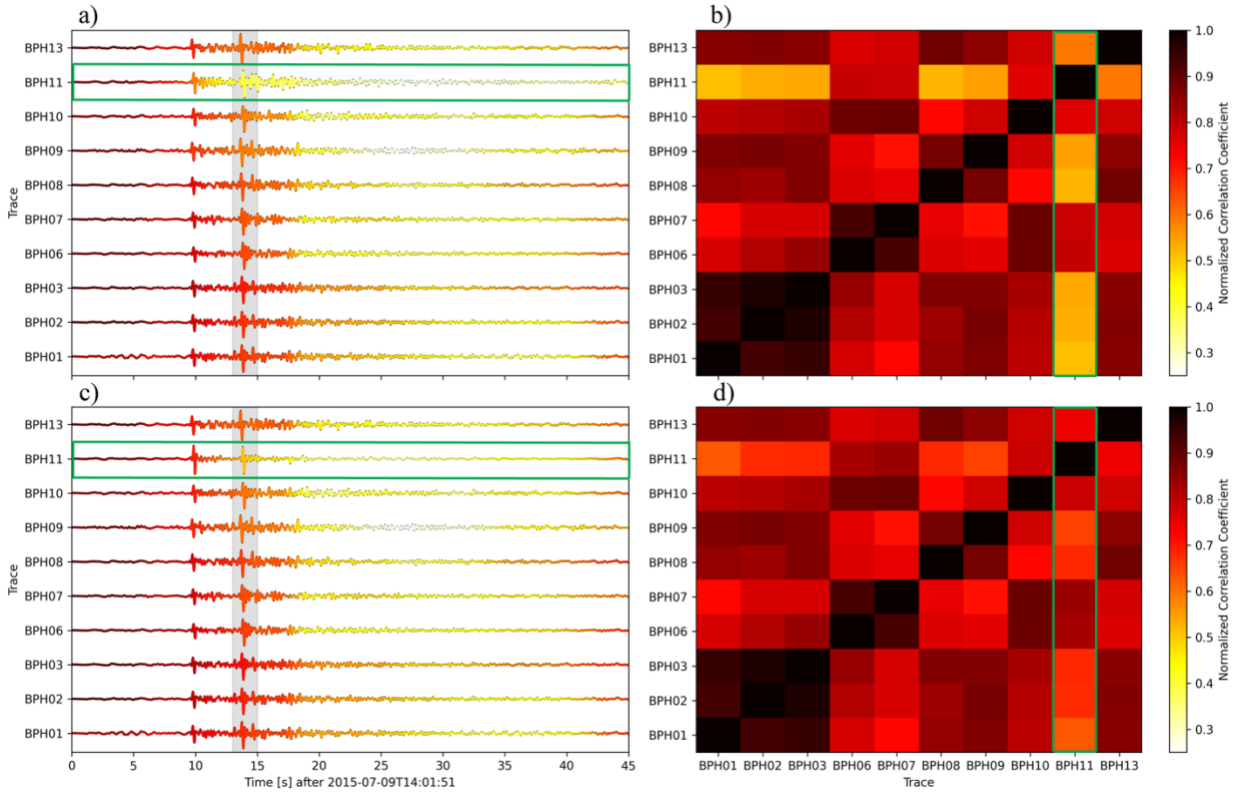


Figure S8: Same description as Figure S7 but now focused on the *S*-wave arrival of another candidate event.

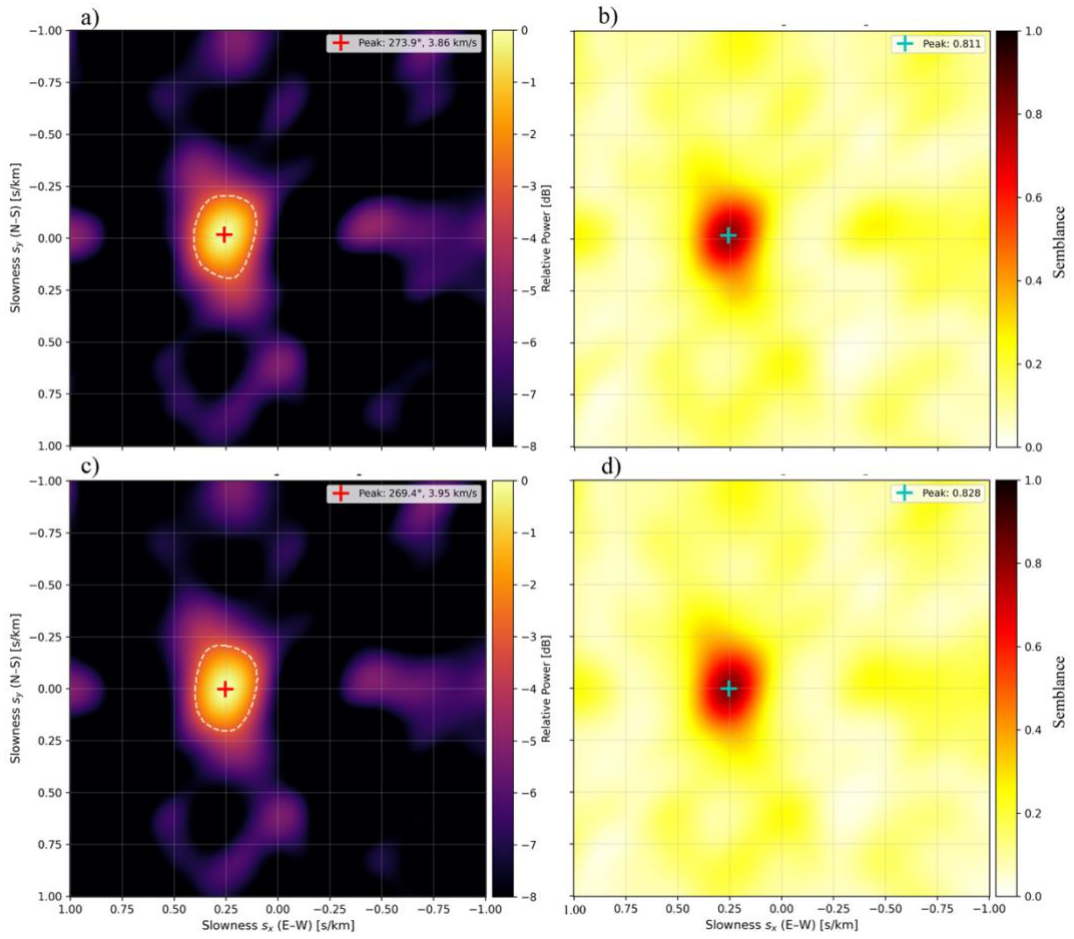


Figure S9: Analysis of the *S*-wave arrival from Figure S8 using *fk* processing with a 2-second window centered on the phase and bandpass filtered between 0.5 – 5 Hz. (a) Relative power distribution for observed waveforms with the peak slowness vector (red cross) and -3 dB contour (dashed). (b) Corresponding semblance, representing array coherence. (c–d) Same as (a–b) for reconstructed waveforms at the masked sensor (BPH11). The reconstructed data recover a consistent slowness vector (269.4°) relative to the observed estimate (273.9°) and closer to the ground-truth back azimuth (264.3°).

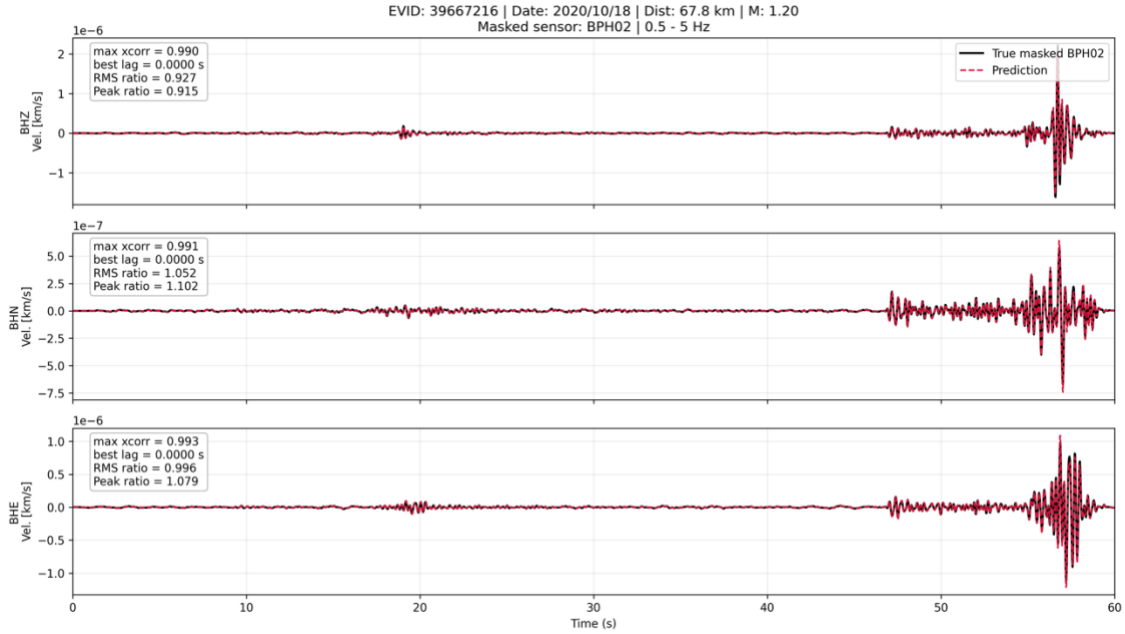


Figure S10: Same as Figure S5, but for an example containing multiple overlapping events within a single window (primary event at ~10 s and a secondary event after ~45 s). The model accurately reconstructs both arrivals, demonstrating robustness to multi-event scenarios.

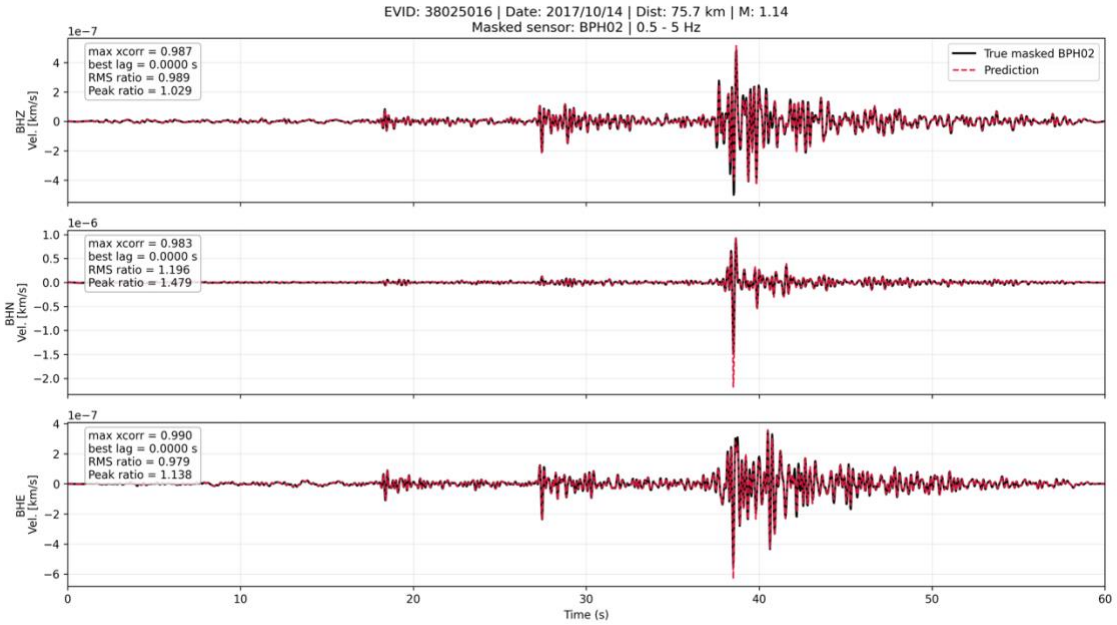


Figure S11: Same as Figure S5, but for an example in which the reconstructed waveform exhibits a temporal offset in the predicted *P*-wave arrival. Despite this shift, the model accurately reconstructs the full waveform with high fidelity, indicating robustness to temporal variability and a lack of strong positional bias.

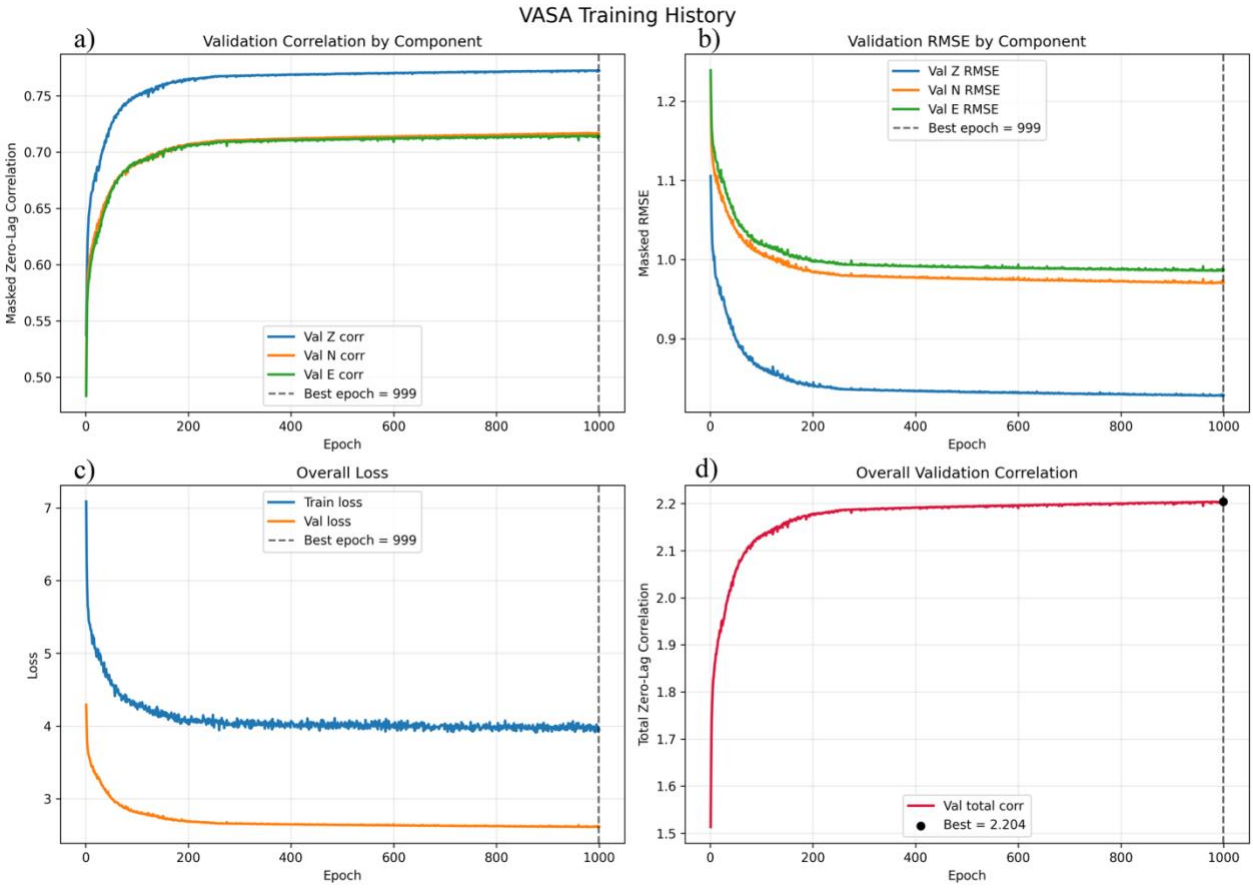


Figure S12: Same as Figure S2, but for the 0.5 – 10 Hz train/test dataset.

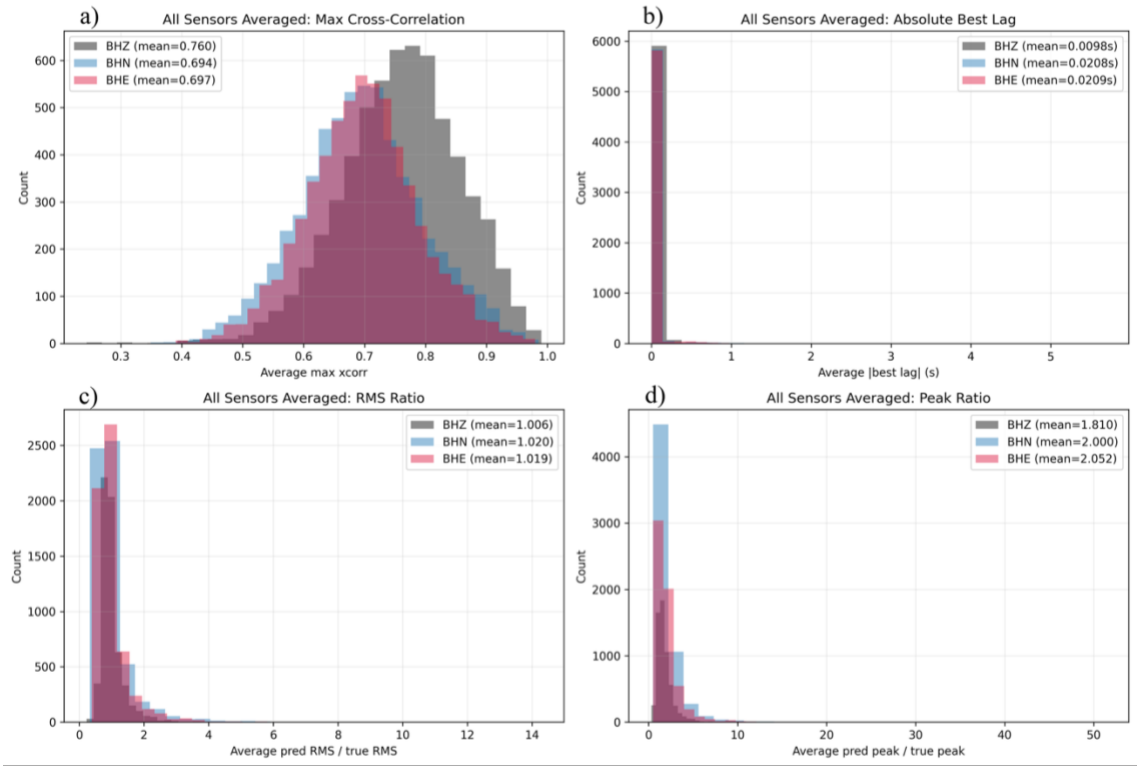


Figure S13: Same as Figure S3, but for the 0.5 – 10 Hz train/test dataset.

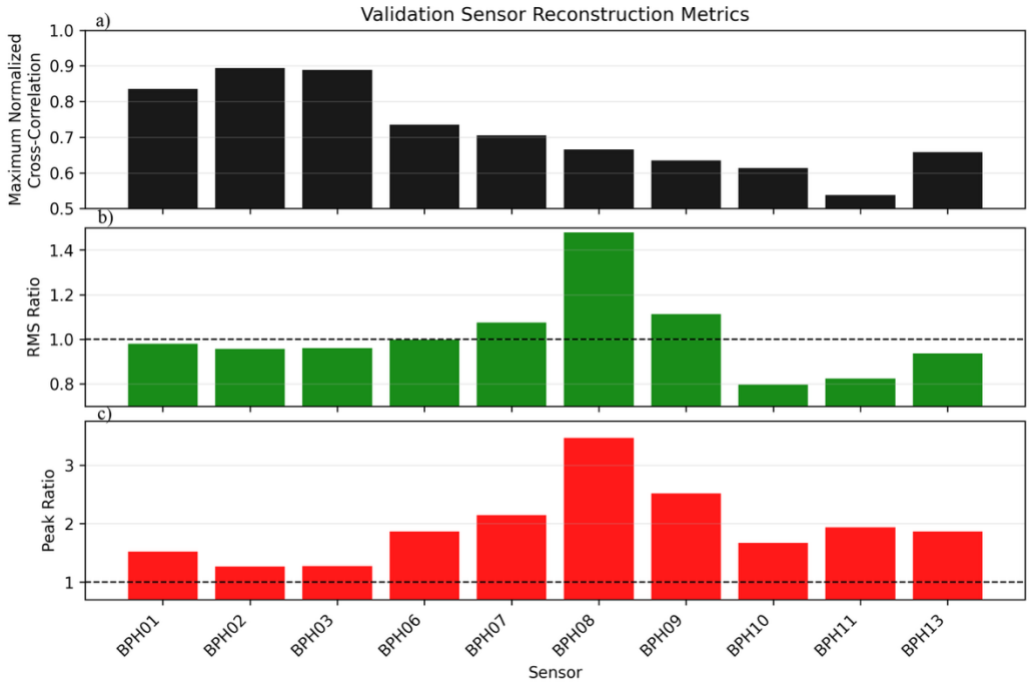


Figure S14: Same as Figure S4, but for the 0.5 – 10 Hz train/test dataset.