

AttentionBased Reconstruction of FullField Tsunami Waves From Sparse Tsunameter Networks

Original

AttentionBased Reconstruction of FullField Tsunami Waves From Sparse Tsunameter Networks / Mcdugald, Edward; Mohan, Arvind; Engwirda, Darren; Marcato, Agnese; E. Santos, Javier. - In: GEOPHYSICAL RESEARCH LETTERS. - ISSN 0094-8276. - 52:16(2025), pp. 1-11. [10.1029/2025gl115345]

Availability:

This version is available at: 11583/3010681 since: 2026-05-08T11:10:47Z

Publisher:

John Wiley and Sons

Published

DOI:10.1029/2025gl115345

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Geophysical Research Letters[®]



RESEARCH LETTER

10.1029/2025GL115345

E.M. thanks the Applied Machine Learning (AML) summer program and Center for Nonlinear Studies (CNLS) at Los Alamos National Laboratory (LANL) for the mentoring and the fellowship received.

Key Points:

- An attention-based neural network specialized for sparse sensing is implemented to reconstruct tsunami waves from sparse observations
- We test our model on tsunami simulations whose epicenters differ from those used in training, replicating realistic forecasting conditions
- We compare our method to the Linear Interpolation with Huygens-Fresnel Principle to generate virtual waveforms, showing improved accuracy

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

E. McDugald,
emcdugald@arizona.edu

Citation:

McDugald, E., Mohan, A., Engwirda, D., Marcato, A., & Santos, J. E. (2025). Attention-based reconstruction of full-field tsunami waves from sparse tsunameter networks. *Geophysical Research Letters*, 52, e2025GL115345. <https://doi.org/10.1029/2025GL115345>

Received 13 FEB 2025
Accepted 1 AUG 2025

Author Contributions:

Conceptualization: Arvind Mohan, Darren Engwirda, Javier E. Santos
Methodology: Edward McDugald, Arvind Mohan, Javier E. Santos
Project administration: Arvind Mohan
Software: Edward McDugald, Darren Engwirda, Agnese Marcato, Javier E. Santos
Supervision: Arvind Mohan, Darren Engwirda, Javier E. Santos

© 2025. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Attention-Based Reconstruction of Full-Field Tsunami Waves From Sparse Tsunameter Networks

Edward McDugald^{1,2} , Arvind Mohan², Darren Engwirda^{3,4} , Agnese Marcato^{5,6} , and Javier E. Santos⁵ 

¹University of Arizona, Tucson, AZ, USA, ²Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, NM, USA, ³Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA, ⁴Environment, Commonwealth Scientific & Industrial Research Organisation, Canberra, NSW, Australia, ⁵Earth and Environmental Sciences, Los Alamos National Laboratory, Los Alamos, NM, USA, ⁶Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM, USA

Abstract We investigate the potential of an attention-based neural network architecture, the Senseiver, for sparse sensing in tsunami forecasting. Specifically, we focus on the Tsunami Data Assimilation Method, which generates forecasts from tsunameter networks. Our model is used to reconstruct high-resolution tsunami wavefields from extremely sparse observations, including cases where the tsunami epicenters are not represented in the training set. Furthermore, we demonstrate that our approach significantly outperforms the Linear Interpolation with Huygens-Fresnel Principle in generating dense observation networks, achieving markedly improved accuracy.

Plain Language Summary While machine learning (ML) methods have achieved accurate forecasts of tsunami waveforms at fixed observation points, the full-field reconstruction of tsunami waves from sparse observations has not yet been demonstrated using ML techniques. This challenge is highly relevant to tsunami data assimilation, where in-ocean tsunameter observations are integrated with numerical models to improve tsunami forecasts. In this study, we employ an attention-based neural network architecture, the Senseiver, to produce high-fidelity reconstructions of tsunami wavefields from extremely sparse tsunameter measurements at realistic sensor locations. To illustrate the practical value of our approach for tsunami forecasters, we perform experiments in which the Senseiver generates virtual waveforms within a sparse observation network, and we compare its performance to a widely used interpolation method for the same task.

1. Introduction

Tsunami early warning systems are critical for mitigating the devastating human and economic impacts of these natural hazards. Most operational systems rely on partial differential equation (PDE)-based solvers, where the initial ocean surface displacement is inferred from earthquake source parameters and then propagated using shallow water equation models. Widely adopted frameworks include NOAA's MOST system (V. Titov & Synolakis, 2016; V. V. Titov & Gonzalez, 1997), GeoClaw (Berger et al., 2011; LeVeque et al., 2011), and Gerris/Basilisk (Popinet, 2012, 2020). Despite broad use, these approaches face significant challenges, including uncertainty in seismic parameter estimation, high computational demand, and limited real-time adaptability.

To overcome these limitations and enable rapid response, tsunami forecasting increasingly leverages sparse offshore and onshore observations for source characterization and direct wave prediction. Methods such as DART buoy data inversion (Percival et al., 2011), direct energy estimation from deep-ocean pressure measurements (Bernard et al., 2013), and real-time data assimilation using GNSS and other sensors (Tsushima & Ohta, 2014) have improved speed and accuracy, yet perform best when observation networks are sufficiently dense and well-placed. This has motivated the development of sparse sensing strategies to enhance observational coverage through techniques such as the generation of virtual waveform observations and optimization of sensor placement. For example, Wang et al.'s (Wang et al., 2019) Huygens-Fresnel interpolation synthesizes virtual tsunami waveforms from sparse tsunameter observations, effectively constructing a denser network that can then be used as input to data assimilation schemes, such as optimal interpolation (Maeda et al., 2015). Concurrently, sensor placement optimization has become an active research area, with Fujita et al. (2024) combining sensor placement and proper orthogonal decomposition for pseudo-super-resolution, and Wang et al. (2020) applying empirical orthogonal function analysis to optimally deploy gauges around Crete Island for maximal early warning

Writing – original draft:

Edward McDugald

Writing – review & editing:

Arvind Mohan, Darren Engwirda,
Agnese Marcato, Javier E. Santos

capability. Notably, these methods are rooted primarily in statistical or physics-based frameworks and have rarely incorporated machine learning (ML) architectures for wavefield reconstruction.

Beyond forecasting arrival times and amplitudes at fixed locations, there is increasing recognition that capturing the spatial dynamics of tsunami waves over wide areas is essential. This not only improves hazard understanding but also supports emergency response, a need highlighted in post-disaster damage assessments using remote sensing and image analysis techniques (Koshimura et al., 2020).

Recent ML approaches show promise in forecasting tsunami waveforms at fixed points (Liu et al., 2021; Rim et al., 2022), but ML solutions for reconstructing full tsunami wavefields from sparse data remain underdeveloped. A related advancement by Archambault et al. (2024) used an attention-based encoder-decoder model to reconstruct sea surface height fields from incomplete satellite altimetry data, leveraging auxiliary sea surface temperature observations and a supervised-to-unsupervised fine-tuning strategy. While their architecture is similar to ours, their problem focuses on spatial interpolation of gridded satellite data, rather than sparse, point-based tsunami sensor networks.

Here, we adapt Senseiver (Santos, 2023), an attention-based neural network tailored for sparse sensing, capable of reconstructing high-resolution tsunami wavefields from limited tsunameter observations. To focus on the fundamental capabilities of this architecture, we adopt an idealized simulation environment that uses parameterized source models rather than fault-derived displacements. We benchmark Senseiver against the Linear Interpolation with Huygens-Fresnel Principle (LIHFP) method (Wang et al., 2019), evaluating accuracy and virtual waveform generation at unobserved locations. Our experiments emulate an operational-like workflow using realistic DART buoy configurations and test generalization to previously unseen epicenters. In addition to generating remarkable full-field reconstructions, results show that Senseiver reduces average errors in wave height, arrival time, and maximum amplitude by approximately a factor of two compared to LIHFP—demonstrating the promise of attention-based ML tools for advancing sensor-informed tsunami forecasting and hazard assessment.

2. Generating Full-Field Training Data

Tsunamis can be effectively modeled using the shallow water equations, which describe the dynamics of a fluid layer via a system of partial differential equations for the depth-averaged horizontal velocity $u = u(\mathbf{x}, t)$ and fluid thickness $h = h(\mathbf{x}, t)$:

$$\frac{\partial u}{\partial t} + (u \cdot \nabla)u + fu^\perp = -(1 - \beta)g\nabla(h + z_b) + c_d \frac{|u|}{h}u + \nu_4 \nabla^4 u, \quad (1)$$

$$\frac{\partial h}{\partial t} + \nabla \cdot (uh) = 0. \quad (2)$$

Here, $z_b = z(\mathbf{x})$ denotes the bathymetric height at the base of the fluid domain, $g = 9.80665 \text{ m s}^{-2}$ is gravitational acceleration, and $f = 2\Omega \sin \phi$ is the Coriolis force, with $\Omega = 7.292 \times 10^{-5} \text{ rad s}^{-1}$ and ϕ the latitude. $u^\perp = \mathbf{k} \times u$ represents tangential velocity due to the Coriolis effect, where \mathbf{k} is the local vertical. The reduced gravity correction $\beta = 0.015$ accounts for self-attraction and loading effects—that is, changes in sea surface height and ocean dynamics resulting from the gravitational pull of redistributed water mass and the elastic deformation of Earth's crust under that load (Inazu & Saito, 2013; Shihora et al., 2022). The coefficients c_d and ν_4 denote drag and viscous scalings, respectively.

Our data set consists of time integrations of these equations, with the initial ocean surface displacement parameterized by the epicenter location \mathbf{x}_0 :

$$h_0(\mathbf{x}; \mathbf{x}_0) = -z_b + 5 \exp\left(-\left(250\|\mathbf{x} - \mathbf{x}_0\|^2\right)^4\right),$$

where $\|\cdot\|$ denotes the Euclidean norm. This initial condition produces a localized, flat-topped solitary wave, serving as a simplified yet physically reasonable proxy for tsunami generation due to seafloor deformation. This approach avoids the complexity of explicitly modeling earthquake fault slip—a technically demanding problem

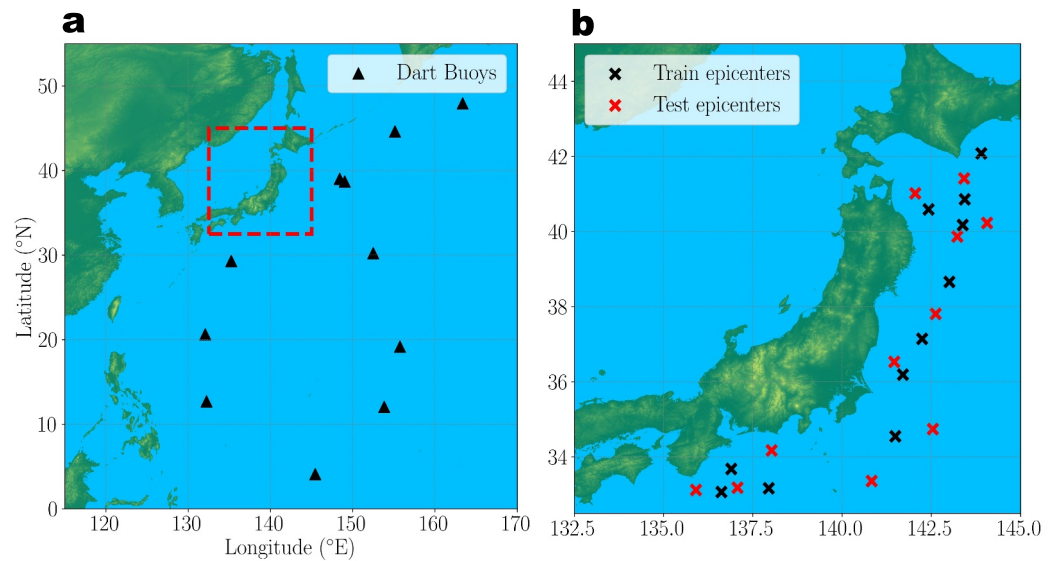


Figure 1. (a) Coverage of DART Buoys, denoted by black triangles. (b) Training (black) and test (red) epicenters used for synthetic tsunami generation.

and an active area of research (Abbate et al., 2024). While more detailed source models are possible (e.g., elastic dislocation or finite fault modeling), adopting a parameterization based solely on epicenter location enables simplicity of experiment design and subsequent analysis. The 5-m displacement represents a significant yet plausible event, consistent with historically observed tsunami amplitudes. For details on the numerical methods, including calibration and validation of the PDE model using historical DART data, see Texts S1–S2 and Figures S1–S2 in Supporting Information S1.

Importantly, this work is structured as a proof of concept for the Senseiver sparse sensing framework. Our primary goal is to evaluate whether Senseiver can reconstruct high-resolution tsunami wavefields from sparse observations, using the deep-ocean DART system as a realistic baseline. DART buoys were selected specifically because they span broad offshore regions, enabling us to assess Senseiver's capacity to infer wavefields over large spatial domains—potentially enhancing direct forecasting and improving data assimilation schemes through virtual waveform augmentation, as demonstrated in (Wang et al., 2019).

To construct our training and test data sets, we queried the USGS earthquake catalog (U.S. Geological Survey, 2023), filtering for events that (a) had magnitudes ≥ 7.5 , (b) listed “Japan” in the location descriptor, and (c) occurred in water depths greater than 1,000 m. The depth constraint was adopted to avoid simulation artifacts from initial conditions placed too close to the coast, which can introduce artifacts on coarser meshes. This also ensures the events used are in operationally relevant offshore regions where DART buoys can realistically contribute data. From an initial pool of 92 qualifying events meeting our criteria, we randomly selected 11 epicenters for training. Test epicenters were drawn from the same pool, ensuring each lies within 20–100 miles of a training source, to emulate scenarios where events are seismologically related but not spatially redundant. This strategy achieves a balance between data diversity and realistic generalization: the model is exposed during training to a broad distribution of real-world offshore earthquakes, while inference is performed on novel but physically consistent examples. All selected epicenters (see Figure 1) are located between 136°–145° longitude and 33°–43° latitude. The test events lie 20–100 miles from the nearest training epicenter, emulating unseen but seismologically relevant cases. A full pairwise distance matrix is provided in Table S1 in Supporting Information S1.

We acknowledge that for typical operational tsunami forecasting near Japan—especially for events generated close to the coast—DART buoy data alone may not be sufficient due to short arrival times. In practice, such forecasts often rely on coastal instrumentation and denser observation networks. Nevertheless, our focus is not on coastal arrival prediction in isolation, but on evaluating the viability of Senseiver in general settings where sparse observations are the primary inputs. We believe Senseiver's demonstrated success across a realistic DART

configuration strongly suggests it would also perform well in other settings, including denser local arrays and finer model meshes.

Each simulation frame consists of 163,842 unstructured spatial points (longitude, latitude, wave height). To reduce computational load, we subsample frames by a factor of two, resulting in 81,921 pixels per time step. Simulations span 4 hr, computed at 50-s intervals (289 time steps), yielding a data set of size (3179, 81921) for each of the training and test sets. Temporal resolution is chosen to balance signal fidelity with computational cost and supports second-order finite difference diagnostics in model validation (see Text S3 in Supporting Information S1).

3. Methods

3.1. The Senseiver

The sparse sensing problem for tsunami waves can be formulated as follows: Given a discrete set of ocean state observations $s = \{s_1, s_2, \dots, s_N\}$ collected at locations $\mathbf{x}^s = \{\mathbf{x}_1^s, \mathbf{x}_2^s, \dots, \mathbf{x}_N^s\}$, the objective is to reconstruct the state $\hat{s} = \{\hat{s}_1, \hat{s}_2, \dots, \hat{s}_M\}$ at a set of query points $\mathbf{x}^q = \{\mathbf{x}_1^q, \mathbf{x}_2^q, \dots, \mathbf{x}_M^q\}$ where typically $M \gg N$. In our work, the input state $s(\mathbf{x}^s)$ corresponds to ocean surface heights measured at DART buoy locations, while the query state $\hat{s}(\mathbf{x}^q)$ represents the ocean surface height throughout the global ocean domain. Since the DART network consists of only a few dozen buoys, yet even coarse global coverage requires hundreds of thousands of points, this reconstruction task is extremely sparse.

To address this challenge, we employ the Senseiver (Santos et al., 2023), an attention-based neural network specifically designed for sparse-to-dense reconstruction tasks. The Senseiver utilizes a multi-level encoder-decoder architecture that leverages the strengths of attention mechanisms and data compression. Key advantages of the attention operation for sparse sensing include: (a) treating positional information as a feature, (b) being agnostic to grid structure, thus accommodating arbitrary meshes, and (c) enabling immediate long-range spatial associations—unlike CNNs, which require deep layers to integrate distant features. However, attention operations scale quadratically with input size. Senseiver overcomes this bottleneck by encoding the input into compact latent arrays, where most attention computations are performed.

The Senseiver workflow is as follows:

1. A positional encoder P_E maps observation locations to spatial encoding vectors \mathbf{a}^s .
2. An attention-based encoder E transforms the observation-value/location pairs (s_i, \mathbf{a}_i^s) into a latent matrix \mathbf{Z} .
3. The positional encoder P_E also maps query locations to spatial encodings \mathbf{a}^q .
4. An attention-based decoder D reconstructs outputs at the encoded query locations \mathbf{a}^q .

The positional encoder P_E implements a trigonometric encoding of spatial coordinates, while the encoder and decoder (E, D) incorporate trainable multi-layer perceptrons within their attention blocks. For a comprehensive description of the Senseiver architecture, see (Santos et al., 2023). The process can be summarized as:

$$\mathbf{a}^s = P_E(\mathbf{x}^s), \quad (3)$$

$$\mathbf{Z} = E(s, \mathbf{a}^s), \quad (4)$$

$$\mathbf{a}^q = P_E(\mathbf{x}^q), \quad (5)$$

$$\hat{s}(\mathbf{x}^q) = D(\mathbf{Z}, \mathbf{a}^q). \quad (6)$$

This flexible architecture enables modeling on unstructured data and allows for custom coordinate choices in the positional encoder. In our application, we augment latitude and longitude with ocean bathymetry encodings. Land pixels are masked out, ensuring training occurs only on ocean data. A schematic of the Senseiver workflow is presented in Figure 2.

Model weights are optimized by minimizing the mean squared error:

$$\mathcal{L} = \sum (s(\mathbf{x}^q) - \hat{s}(\mathbf{x}^q))^2, \quad (7)$$

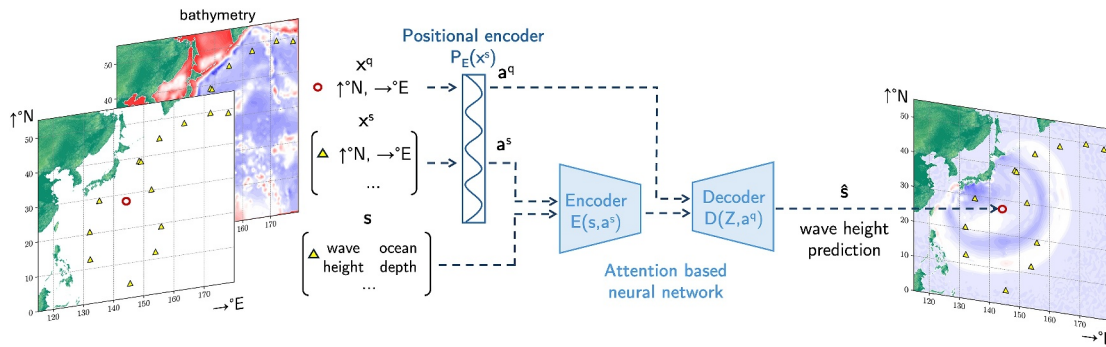


Figure 2. Overview of the Senseiver architecture for tsunami reconstruction. The sensor values along with their encoded positions are processed by the encoder into a latent array Z of fixed size. The encoded query locations are processed by the attention-based decoder, which outputs an estimate of wave height at the query location(s).

$$\text{where } \hat{s}(\mathbf{x}^q) = D(E(s(\mathbf{x}^q), P_E(\mathbf{x}^s)), P_E(\mathbf{x}^q)), \quad (8)$$

using the Adam optimizer (Kingma & Ba, 2014). To prevent overfitting, we train on only 80% of the available training frames. The number and ordering of query points, as well as the data frames used in training, are treated as hyperparameters, with further details provided in Text S4 in Supporting Information S1.

3.2. Experiments

The primary error metric used to evaluate reconstruction quality is the mean absolute error (MAE), normalized by the absolute maximum wave height in each frame. To ensure this metric captures meaningful differences, we compute the MAE only over pixels where the true wave height exceeds 1×10^{-4} . This threshold is necessary because most of the ocean domain has zero wave height; including these pixels would artificially deflate the error. For each reconstruction, we calculate the absolute difference between the true and reconstructed wave height fields, divided by the maximum absolute true wave height for that frame:

$$\text{Error}(x, y, t) = \frac{|h(x, y, t) - \hat{h}(x, y, t)|}{\max(|h(x, y, t)|)}. \quad (9)$$

While this metric provides a standard quantitative assessment typical in ML, we also sought to evaluate our approach in a more operationally relevant context. In practice, tsunami forecasting systems such as the tsunami data assimilation method (Maeda et al., 2015) focus on predicting key features of tsunami waveforms—specifically, arrival times and maximum amplitudes—at fixed observation points. TDAM performs well with dense observation networks, but its accuracy diminishes with sparse sensor coverage. To address this, the Linear Interpolation with Huygens-Fresnel Principle (LIHFP) (Wang et al., 2019) was developed to augment sparse networks, estimating virtual waveforms at unsensed locations by interpolating between real sensor data while accounting for arrival times, distances, and bathymetry.

To bridge the gap between standard ML evaluation and practical tsunami forecasting, we compare our model's reconstructions not only in terms of full-field MAE, but also by generating six virtual waveforms at strategically chosen locations and benchmarking them against those produced by the LIHFP method. This dual evaluation strategy allows us to assess both the general reconstruction capability of our model and its effectiveness in a real-world forecasting workflow.

4. Results

Averaging the reconstruction loss (Equation 9) over all 289 frames from each of the 11 test epicenters yields a global experiment error of 5.00×10^{-2} . Notably, much of this error arises from the early time frames, when sensors have yet to detect significant tsunami signals. This limitation is inherent to any assimilation method that relies solely on tsunameter data. On average, the reconstruction error remains below 1×10^{-1} after 32.9 min, with a median threshold time of 38.3 min. To assess the physical consistency of our model's reconstructions, we

evaluated the continuity equation residual (Equation 2), yielding a global average error of 4.10×10^{-2} across all times and epicenters. Details of the diagnostic calculation and frame-by-frame residual plots are provided in the Supporting Information (Text S3 and Figures S15 in Supporting Information S1). The key results from the reconstruction experiments for all eight simulations are summarized in the first seven columns of Table 1.

Figure 3 presents representative reconstructions for two test simulations, with epicenters at ($143.4^{\circ}E, 41.4^{\circ}N$) and ($140.8^{\circ}E, 33.4^{\circ}N$), shown at 75, 150, and 225 min. For each case, we include example waveform reconstructions at fixed DART buoy locations and the per-pixel average reconstruction error over time. These two epicenters were chosen to illustrate model performance at different distances from the training set: the first is the median-distance epicenter (38.4 miles from the nearest training example), while the second is the most distant (90.6 miles). Together, they demonstrate reconstruction accuracy for both typical and challenging test cases. Reconstructions, error plots, and time series for all other test epicenters are provided in Figures S3–S14 in Supporting Information S1.

Visual inspection shows that the Senseiver effectively captures key features of the tsunami wavefield, including the spatial support, number of wave periods, and amplitude. The error plots reveal that reconstruction error decreases rapidly within the first ≈ 50 min, highlighting the particular challenge of early time reconstruction when little sensor information is available. For the median-distance epicenter, the error drops below 1×10^{-1} after just 10 min; for the most distant epicenter, this threshold is reached at 40.8 min. While resolving the earliest time frames is not the focus of this work, we suggest alternative data sets that could better inform the Senseiver at early time frames in the conclusion.

To ensure consistency with the experimental setup of (Wang et al., 2019) and enable direct comparison, we focused our evaluation on virtual waveform reconstruction at midpoints between selected DART buoys. Although our experiments center on these midpoints, the methodology is readily extensible to arbitrary virtual sensor locations. In fact, analyzing Senseiver performance as the virtual observation point moves farther from the training data may provide valuable insight into the method's effective range—potentially revealing where LIHFP waveforms surpass Senseiver reconstructions in accuracy. For this study, we identified six midpoints among DART buoys that registered significant waveforms in our simulations and compared the Senseiver's reconstructions at these locations to those produced by LIHFP.

Across the majority of test epicenters, times, and virtual observation sites, the Senseiver outperformed LIHFP, achieving a MAE in wave height of 7.02×10^{-2} meters, compared to 18.2×10^{-2} meters for LIHFP. The Senseiver also demonstrated substantial improvements in estimating both arrival times and maximum amplitudes at nearly all virtual observation points. Specifically, the MAE for arrival time, averaged across all 11 simulations, was 8.12 min for the Senseiver and 16.68 min for LIHFP. The only exception occurred at the epicenter located at ($135.9^{\circ}E, 33.1^{\circ}N$), where the Senseiver's arrival time MAE was 18.611 min, slightly higher than the 17.778 min observed for LIHFP. For maximum amplitude, the Senseiver achieved an MAE of 13.1×10^{-2} meters, outperforming LIHFP's MAE of 25.9×10^{-2} meters across all cases. On average, the percent change in arrival time MAE for the Senseiver was -51.9% , and for maximum amplitude, it was -49.9% , indicating a substantial reduction in error compared to LIHFP. These results are summarized in Figures 4b and 4c, which display arrival time and maximum amplitude errors for both methods across all 48 test epicenter–virtual observation point pairings. Key numerical results for all 11 simulations can be found in the last six columns of Table 1.

Figure 4d provides a detailed summary for the epicenter at ($143.4^{\circ}E, 41.4^{\circ}N$) (the median epicenter). For this case, the Senseiver achieved a substantially lower MAE across the six virtual waveforms: 5.5×10^{-2} meters, compared to 13.3×10^{-2} meters for LIHFP. Arrival time errors were similarly reduced, with the Senseiver yielding an MAE of 7.22 versus 16.39 min for LIHFP. For maximum amplitude, the Senseiver achieved an MAE of 14.9×10^{-2} meters, again outperforming LIHFP's 23.8×10^{-2} meters. Virtual waveform comparisons for all remaining test epicenters are provided in Figures S16–S24 in Supporting Information S1.

We conclude this section by noting the lack of temporal smoothness observed in the waveforms generated by the Senseiver. To mitigate this issue, we applied median filtering—a non-linear technique that replaces each point in a signal with the median value within a local window, effectively reducing noise and outliers while preserving key features—in our virtual waveform experiments. This post-processing step substantially improved arrival time accuracy: for arrival times defined by a 3 cm threshold, repeated median filtering reduced the raw Senseiver error

Table 1
Overview of Senseiver Evaluation Metrics for All Test Epicenters

Epicenter number	Longitude	Latitude	Closest distance from training set	Mean recon error	Mean physics error	Trigger time (mins)
1	140.827	33.362	90.551	0.073	0.047	40.8
2	141.459	36.534	27.105	0.045	0.043	38.3
3	142.050	41.019	35.398	0.056	0.045	46.7
4	142.542	34.745	61.320	0.062	0.040	43.3
5	143.228	39.869	22.648	0.036	0.036	14.2
6	144.060	40.232	36.197	0.042	0.039	39.2
7	142.619	37.812	50.405	0.040	0.035	24.2
8	143.416	41.415	38.360	0.041	0.045	10.0
9	135.905	33.123	41.433	0.061	0.043	70.8
10	137.071	33.184	27.370	0.035	0.036	11.7
11	138.025	34.175	69.802	0.057	0.041	22.5

Senseiver arrival time MAE (mins)	LIHFP arrival time MAE (mins)	Senseiver max amp MAE (m)	LIHFP max amp MAE (m)	Senseiver wave height MAE (m)	LIHFP wave height MAE (m)
13.333	17.361	0.100	0.326	0.120	0.246
8.333	16.458	0.121	0.316	0.064	0.211
5.972	15.833	0.067	0.261	0.054	0.117
8.611	17.361	0.229	0.363	0.103	0.268
6.111	15.417	0.186	0.258	0.054	0.160
5.000	15.903	0.185	0.248	0.069	0.167
3.750	16.181	0.182	0.291	0.073	0.201
7.222	16.389	0.149	0.238	0.055	0.133
18.611	17.778	0.100	0.187	0.087	0.165
4.722	17.431	0.043	0.161	0.047	0.173
7.639	17.361	0.082	0.198	0.068	0.164

Note. Columns 4–7 correspond to reconstruction experiments, and columns 8–13 correspond to LIHFP comparisons. MAE refers to mean absolute error, and Trigger Time refers to the time at which no reconstruction exceeds 1×10^{-1} .

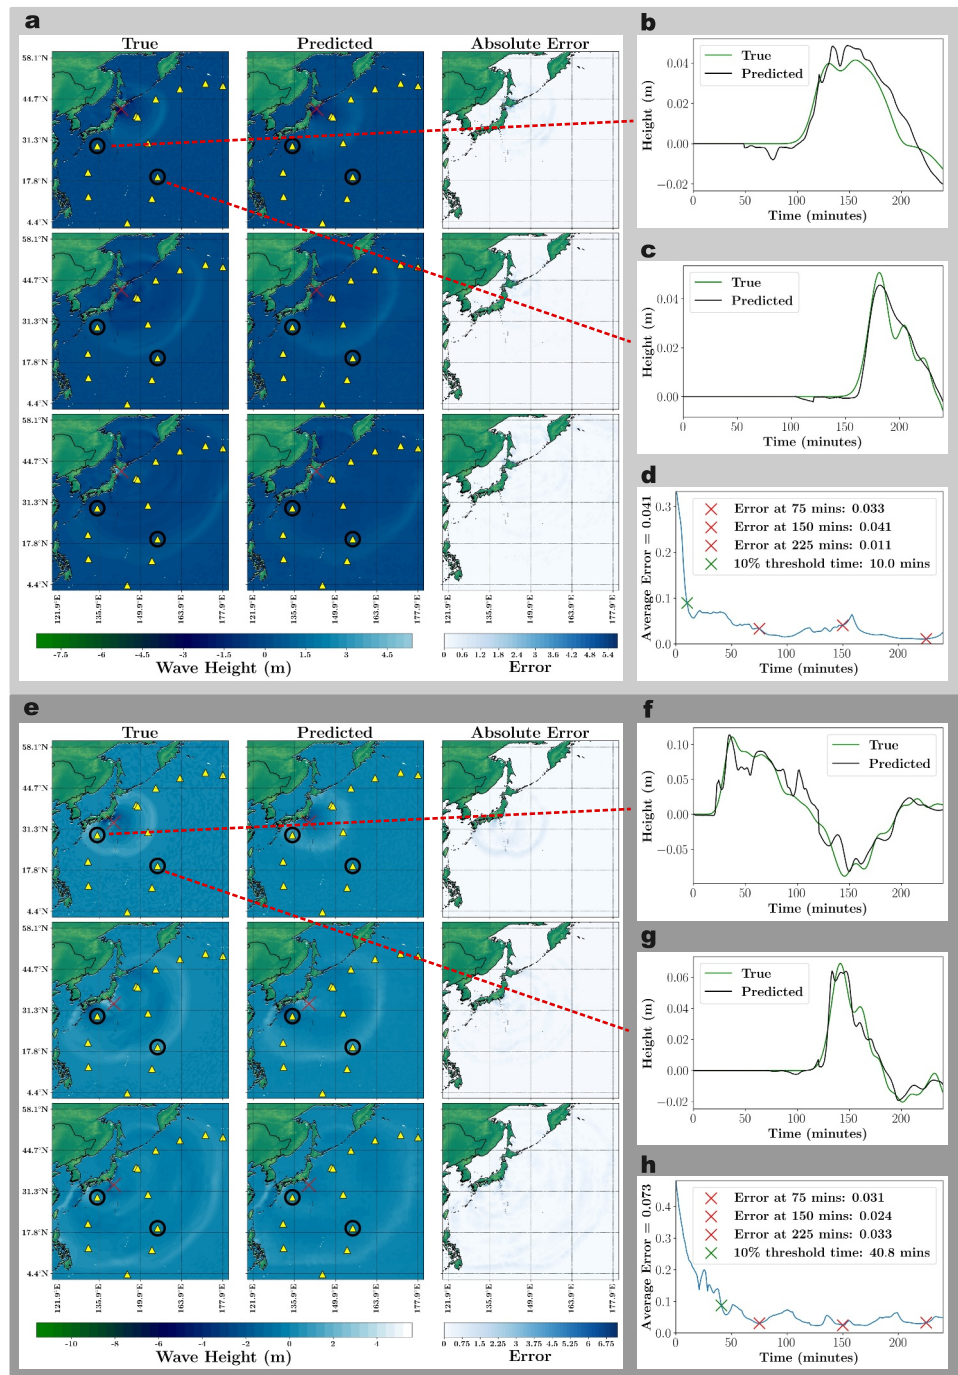


Figure 3. (a) Full reconstruction at 75, 150, and 225 min for test epicenter ($143.4^{\circ}E, 41.4^{\circ}N$). The epicenter location is indicated by an X, and yellow triangles denote the sensors. (b)–(c) Waveform reconstructions at fixed DART buoys with longitude-latitude coordinate ($135.2^{\circ}E, 29.3^{\circ}N$) and ($155.7^{\circ}E, 19.2^{\circ}N$), respectively from top to bottom. (d) Per-pixel average reconstruction error as a function of time for test epicenter ($143.4^{\circ}E, 41.4^{\circ}N$). (e)–(h) Equivalent figures for test epicenter ($140.8^{\circ}E, 33.4^{\circ}N$).

from 9.05 to 5.35 min. However, this improvement came at the cost of a decrease in maximum amplitude accuracy. For the plots in Figure 4 and the statistics collected in Table 1, Senseiver signals were processed using a single median filter with a kernel length of 13. Analogous statistics are provided for the raw Senseiver waveforms as well as heavily filtered senseiver waveforms in Tables S2–S3 in Supporting Information S1. Addressing temporal smoothness directly within the ML framework—for example, by incorporating temporal encodings,

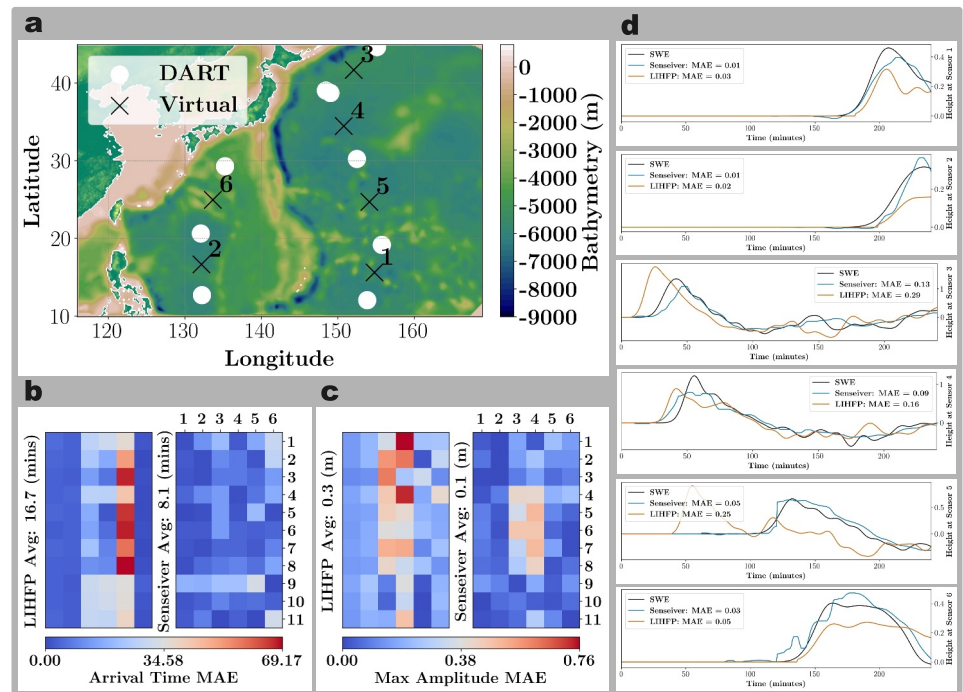


Figure 4. (a) Observation network augmented with virtual sensors with locations derived as midpoints between DART buoys. (b)–(c). Mean absolute error in arrival time and maximum amplitude, respectively. Each color map (6 columns \times 11 rows) represents all combinations of test epicenter and virtual observation pairs. (d) Senseiver and LIHFP-based virtual waveform reconstructions at test epicenter ($143.4^\circ E, 41.4^\circ N$).

conditioning on previous states, or introducing regularization terms into the loss function—remains an important direction for future research.

5. Conclusion

We have presented Senseiver, an attention-based neural network capable of reconstructing high-resolution tsunami wavefields from sparse observation networks. Our results demonstrate substantial improvements over the LIHFP baseline, particularly in generating accurate virtual waveforms at unobserved locations. To our knowledge, this is the first ML-based solution to address the sparse sensing problem for tsunami forecasting at high spatial resolution.

This work serves as a proof of concept, with several important limitations. Our experiments are restricted to a limited oceanic region and idealized initial conditions. Future research should extend these methods to broader domains, diversify training data, and investigate how regional scale interacts with model complexity. Incorporating additional features and regularization strategies may further improve the temporal smoothness of Senseiver-generated waveforms.

A key challenge remains early stage reconstruction, before tsunameter data are available. Future work could explore integrating alternative data sources such as GNSS (Rim et al., 2022), HF radar (Wang et al., 2023), or distributed acoustic sensing (Xiao et al., 2024) to provide earlier or more robust wavefield estimates. Additionally, adapting Senseiver to optimize sensor placement during training (Marcato et al., 2024) could inform the design of observation networks, particularly in resource-limited regions.

By bridging advances in ML with operational tsunami sensing, Senseiver offers a promising step toward more accurate and timely tsunami early warning systems.

Data Availability Statement

The tsunami simulation data used in this study are available from Zenodo <https://zenodo.org/records/15478821> (McDugald, 2025). The Senseiver codebase is available from GitHub <https://github.com/OrchardLANL/Senseiver> (Santos, 2023), with tsunami-specific code provided in the “tsunami” directory. Formal citations for both the data set and software are included in the reference list.

Acknowledgments

This project was partially funded by the Artemis LDRD program at the Los Alamos National Laboratory. The Los Alamos Unlimited Release number for this work is LA-UR-24-32151. Andrew Arnold (University of Arizona PhD Student in Applied Mathematics) assisted in analyzing the virtual waveform experiments. Additional references used only in the Supporting Information S1 include: Barton et al. (2022), Engwirda (2017), Fuji et al. (2011), GEBCO Compilation Group (2023), Lilly, Capodaglio, et al. (2023), Lilly, Engwirda, et al. (2023), NOAA (2011), Pal et al. (2023), Ringler et al. (2010), and Thuburn et al. (2009).

References

- Abbate, A., González Vida, J. M., Castro Díaz, M. J., Romano, F., Bayraktar, H. B., Babeyko, A., & Lorito, S. (2024). Modelling tsunami initial conditions due to rapid coseismic seafloor displacement: Efficient numerical integration and a tool to build unit source databases. *Natural Hazards and Earth System Sciences*, 24(8), 2773–2791. <https://doi.org/10.5194/nhess-24-2773-2024>
- Archambault, T., Filoche, A., Charantonis, A., Béréziat, D., & Thiria, S. (2024). Learning sea surface height interpolation from multi-variate simulated satellite observations. *Journal of Advances in Modeling Earth Systems*, 16(6), e2023MS004047. <https://doi.org/10.1029/2023MS004047>
- Barton, K. N., Pal, N., Brus, S. R., Petersen, M. R., Arbic, B. K., Engwirda, D., et al. (2022). Global barotropic tide modeling using inline self-attraction and loading in MPAS-Ocean. *Journal of Advances in Modeling Earth Systems*, 14(11), e2022MS003207. <https://doi.org/10.1029/2022ms003207>
- Berger, M. J., George, D. L., LeVeque, R. J., & Mandli, K. T. (2011). The GeoClaw software for depth-averaged flows with adaptive refinement. *Advances in Water Resources*, 34(9), 1195–1206. <https://doi.org/10.1016/j.advwatres.2011.02.016>
- Bernard, E., Wei, Y., Tang, L., & Titov, V. (2013). Impact of near-field, deep-ocean tsunami observations on forecasting the 7 December 2012 Japanese tsunami. *Pure and Applied Geophysics*, 171(12), 3483–3491. <https://doi.org/10.1007/s00024-013-0720-8>
- Engwirda, D. (2017). JIGSAW-GEO: Locally orthogonal staggered unstructured grid generation for general circulation modelling on the sphere. *Geoscientific Model Development*, 10(6), 2117–2140. <https://doi.org/10.5194/gmd-10-2117-2017>
- Fujii, Y., Satake, K., Sakai, S., Shinohara, M., & Kanazawa, T. (2011). Tsunami source of the 2011 off the Pacific coast of Tohoku Earthquake. *Earth Planets and Space*, 63(7), 815–820. <https://doi.org/10.5047/eps.2011.06.010>
- Fujita, S., Nomura, R., Moriguchi, S., Otake, Y., Koshimura, S., Leveque, R., & Terada, K. (2024). Optimization of a tsunami gauge configuration for pseudo-super-resolution of wave height distribution. *Earth and Space Science*, 11(2), e2023EA003144. <https://doi.org/10.1029/2023EA003144>
- GEBCO Compilation Group. (2023). GEBCO 2023 grid. <https://doi.org/10.5285/f98b053b-0cbc-6c23-e053-6c86abc0af7b>
- Inazu, D., & Saito, T. (2013). Simulation of distant tsunami propagation with a radial loading deformation effect. *Earth Planets and Space*, 65(8), 835–842. <https://doi.org/10.5047/eps.2013.03.010>
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Koshimura, S., Moya, L., Mas, E., & Bai, Y. (2020). Tsunami damage detection with remote sensing: A review. *Geosciences*, 10(5), 177. <https://doi.org/10.3390/geosciences10050177>
- LeVeque, R. J., George, D. L., & Berger, M. J. (2011). Tsunami modelling with adaptively refined finite volume methods. *Acta Numerica*, 20, 211–289. <https://doi.org/10.1017/s0962492911000043>
- Lilly, J. R., Capodaglio, G., Petersen, M. R., Brus, S. R., Engwirda, D., & Higdon, R. L. (2023). Storm surge modeling as an application of local time-stepping in MPAS-Ocean. *Journal of Advances in Modeling Earth Systems*, 15(1), e2022MS003327. <https://doi.org/10.1029/2022ms003327>
- Lilly, J. R., Engwirda, D., Capodaglio, G., Higdon, R. L., & Petersen, M. R. (2023). CFL optimized forward-backward Runge-Kutta schemes for the shallow-water equations. *Monthly Weather Review*, 151(12), 3191–3208. <https://doi.org/10.1175/mwr-d-23-0113.1>
- Liu, C., Rim, D., Baraldi, R., & Leveque, R. (2021). Comparison of machine learning approaches for tsunami forecasting from sparse observations. *Pure and Applied Geophysics*, 178(12), 5129–5153. <https://doi.org/10.1007/s00024-021-02841-9>
- Maeda, T., Obara, K., Shinohara, M., Kanazawa, T., & Uehira, K. (2015). Successive estimation of a tsunami wavefield without earthquake source data: A data assimilation approach toward real-time tsunami forecasting. *Geophysical Research Letters*, 42(19), 7923–7932. <https://doi.org/10.1002/2015GL065588>
- Marcato, A., Gultinan, E., Viswanathan, H., O'Malley, D., Lubbers, N., & Santos, J. E. (2024). Journey over destination: Dynamic sensor placement enhances generalization. *Machine Learning: Science and Technology*, 5(2), 025070. Retrieved from <https://iopscience.iop.org/article/10.1088/2632-2153/ad4e06/meta>
- McDugald, E. (2025). Tsunami sims v2 [dataset]. *Zenodo*. <https://doi.org/10.5281/zenodo.15478821>
- NOAA. (2011). National Data Buoy Center – DART. Retrieved from <http://www.ndbc.noaa.gov/dart.shtml>
- Pal, N., Barton, K. N., Petersen, M. R., Brus, S. R., Engwirda, D., Arbic, B. K., et al. (2023). Barotropic tides in MPAS-Ocean (E3SM V2): Impact of ice shelf cavities. *Geoscientific Model Development*, 16(4), 1297–1314. <https://doi.org/10.5194/gmd-16-1297-2023>
- Percival, D., Denbo, D., Eble, M., Gica, E., Mofjeld, H., Spillane, M., et al. (2011). Extraction of tsunami source coefficients via inversion of dart circledR buoy data. *Natural Hazards*, 58(1), 567–590. <https://doi.org/10.1007/s11069-010-9688-1>
- Popinet, S. (2012). Adaptive modelling of long-distance wave propagation and fine-scale flooding during the Tohoku tsunami. *Natural Hazards and Earth System Sciences*, 12(4), 1213–1227. <https://doi.org/10.5194/nhess-12-1213-2012>
- Popinet, S. (2020). A vertically-Lagrangian, non-hydrostatic, multilayer model for multiscale free-surface flows. *Journal of Computational Physics*, 418, 109609. <https://doi.org/10.1016/j.jcp.2020.109609>
- Rim, D., Baraldi, R., Liu, C. M., LeVeque, R. J., & Terada, K. (2022). Tsunami early warning from global navigation satellite system data using convolutional neural networks. *Geophysical Research Letters*, 49(20), e2022GL099511. <https://doi.org/10.1029/2022gl099511>
- Ringler, T., Thuburn, J., Klemp, J., & Skamarock, W. (2010). A unified approach to energy conservation and potential vorticity dynamics for arbitrarily-structured C-grids. *Journal of Computational Physics*, 229(9), 3065–3090. <https://doi.org/10.1016/j.jcp.2009.12.007>
- Santos, J. E. (2023). The senseiver codebase [software]. *GitHub*. Retrieved from <https://github.com/OrchardLANL/Senseiver>
- Santos, J. E., Fox, Z. R., Mohan, A., O'Malley, D., Viswanathan, H., & Lubbers, N. (2023). Development of the Senseiver for efficient field reconstruction from sparse observations. *Nature Machine Intelligence*, 5(11), 1–9. <https://doi.org/10.1038/s42256-023-00746-x>
- Shihora, L., Sulzbach, R., Dobszlaw, H., & Thomas, M. (2022). Self-attraction and loading feedback on ocean dynamics in both shallow water equations and primitive equations. *Ocean Modelling*, 169, 101914. <https://doi.org/10.1016/j.ocemod.2021.101914>

- Thuburn, J., Ringler, T., Skamarock, W., & Klemp, J. (2009). Numerical representation of geostrophic modes on arbitrarily structured C-grids. *Journal of Computational Physics*, 228(22), 8321–8335. <https://doi.org/10.1016/j.jcp.2009.08.006>
- Titov, V., K anođlu, U., & Synolakis, C. (2016). Development of most for real-time tsunami forecasting. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 142(6), 03116004. [https://doi.org/10.1061/\(ASCE\)WW.1943-5460.0000357](https://doi.org/10.1061/(ASCE)WW.1943-5460.0000357)
- Titov, V. V., & Gonzalez, F. I. (1997). Implementation and testing of the method of splitting tsunami (MOST) model.
- Tsushima, H., & Ohta, Y. (2014). Review on near-field tsunami forecasting from offshore tsunami data and onshore GNSS data for tsunami early warning. *Journal of Disaster Research*, 9(3), 339–357. <https://doi.org/10.20965/jdr.2014.p0339>
- U.S. Geological Survey. (2023). *Earthquake Catalog Search: U.S. Geological Survey*. Retrieved from <https://earthquake.usgs.gov/earthquakes/search/>
- Wang, Y., Heidarzadeh, M., Satake, K., Mulia, I. E., & Yamada, M. (2020). A tsunami warning system based on offshore bottom pressure gauges and data assimilation for Crete Island in the eastern mediterranean basin. *Journal of Geophysical Research: Solid Earth*, 125(10), e2020JB020293. <https://doi.org/10.1029/2020JB020293>
- Wang, Y., Imai, K., Mulia, I. E., Ariyoshi, K., Takahashi, N., Sasaki, K., et al. (2023). Data assimilation using high-frequency radar for tsunami early warning: A case study of the 2022 Tonga volcanic tsunami. *Journal of Geophysical Research: Solid Earth*, 128(2), e2022JB025153. <https://doi.org/10.1029/2022jb025153>
- Wang, Y., Maeda, T., Satake, K., Heidarzadeh, M., Su, H.-Y., Sheehan, A., & Gusman, A. (2019). Tsunami data assimilation without a dense observation network. *Geophysical Research Letters*, 46(4), 2045–2053. <https://doi.org/10.1029/2018GL080930>
- Xiao, H., Spica, Z., Li, J., & Zhan, Z. (2024). Detection of earthquake infragravity and tsunami waves with underwater distributed acoustic sensing. *Geophysical Research Letters*, 51(2), e2023GL106767. <https://doi.org/10.1029/2023GL106767>