

A Source Clustering Approach for Efficient Inundation Modeling and Regional Scale PTHA

Amy L. Williamson^{1,*}, Donsub Rim², Loyce M. Adams³, Randall J. LeVeque³,
Diego Melgar¹, and Frank I. Gonzalez⁴

¹University of Oregon, Department of Earth Sciences, Eugene, Oregon, USA

²New York University, Courant Institute, New York, New York, USA

³University of Washington, Department of Applied Mathematics, Seattle, Washington, USA

⁴University of Washington, Department of Earth and Space Sciences, Seattle, Washington, USA

Correspondence*:
Amy Williamson
awillia5@uoregon.edu

ABSTRACT

For coastal regions on the margin of a subduction zone, near-field megathrust earthquakes are the source of the most extreme tsunami hazards, and are important to handle properly as one aspect of any Probabilistic Tsunami Hazard Assessment (PTHA). Typically, great variability in inundation depth at any point is possible due to the extreme variation in extent and pattern of slip over the fault surface. This is particularly true for long subduction zones such as Cascadia or Sumatra, which are known to have had both full-length and partial ruptures in the past. Moreover, the inundation depth depends greatly on the degree of coastal uplift or subsidence in the area, which can also vary greatly between events.

In this context, we present an approach to estimating inundation depth probabilities (in the form of hazard curves at a set of coastal locations) that consists of two components. The first component uses a Karhunen-Loève expansion to express the probability density function (PDF) for all possible events, using a method developed in previous work by the authors, with PDF parameters that are geophysically reasonable for the Cascadia Subduction Zone (CSZ). It is then extremely easy and computationally cheap to generate a very large N number of samples from this PDF; doing so and performing a full tsunami inundation simulation for each provides a brute force approach to estimating probabilities of inundation. The estimate is simply the number of random realizations that reach specified inundation depth divided by N . However, to obtain reasonable results, particularly for extreme flooding due to rare events, N would have to be so large as to make the necessary tsunami simulations prohibitively expensive. The second component tackles this difficulty by using importance sampling techniques to ensure we adequately sample the tails of the distribution and properly re-weight the probability assigned to the resulting realizations, and by grouping the realizations into a small number of clusters that we believe will give similar inundation patterns in the region of interest. In this approach, only one fine-grid tsunami simulation need be computed from a representative member of each

cluster, and weighted by the sum of the weights of all realizations in the cluster. We discuss clustering approaches based on proxy quantities that are cheap to compute over a large number of realizations, but that can identify a smaller number of clusters of realizations that will have similar inundation depths. The fine-grid simulations for each cluster representative can also be used to develop an improved strategy, in which these are combined with cheaper coarse-grid simulations of other members of the cluster. We illustrate the methodology by considering two coastal locations on the CSZ margin, Crescent City, CA and Westport, WA.

Keywords: probabilistic tsunami hazard assessment, clustering, stochastic earthquakes, Karhunen-Loève expansion, GeoClaw

1 INTRODUCTION

The primary goal of this work is to present a general methodology for developing the *hazard curve* for a quantity of interest (e.g. maximum water depth) at a coastal location that may be inundated by tsunamis. An inundation hazard curve shows the annual probability that the flooding depth will exceed each value for a range of specified *exceedance values*. The same techniques could be applied to other quantities of interest (e.g., maximum flow speed or momentum flux) but here we concentrate on water depth h for illustration, and use h_{\max} to represent the maximum value of h over the entire simulation (at some particular point of interest). Let \hat{h} denote some particular exceedance value. The hazard curve is then obtained by determining $P[h_{\max} > \hat{h}]$, the *annual* probability that h_{\max} exceeds \hat{h} at this particular location, as a function of \hat{h} . The ultimate goal is to develop such a hazard curve at each point on a fine grid covering a community of interest, from which it is possible to then create *hazard maps* that show the spatial distribution of maximum water depth expected for a given annual probability, or the spatial distribution of annual probability for a given exceedance value, or potentially other products useful to emergency managers or community planners. We give some examples of how this can be done in Section 6.

A full probabilistic tsunami hazard assessment (PTHA) would have to include all potential sources of tsunamis, far-field as well as near-field, and also possibly tsunamis induced by landslides or other processes; a complete review of this can be found in [Grezio et al. \(2017\)](#). Here we concentrate on one aspect of PTHA, assessing the probabilities in a coastal region due to a megathrust event on a nearby subduction zone. This is a difficult aspect of PTHA because variations in the spatial distribution of the slip can have a significant effect on the resulting tsunami ([Goda et al., 2016](#); [Melgar et al., 2019](#)). In addition, some events may cause substantial subsidence or uplift of the coast around the location of interest, which can also greatly effect the inundation extent and depth of the resulting tsunami.

To perform PTHA for such events it is necessary to first have some model for the probability density function (PDF) of all possible events. It is impossible to know the correct distribution due to the high degree of epistemic uncertainty in subduction zones with infrequent past megathrust events. However, recent studies have suggested ways to generate a geophysically reasonable distribution that can be easily sampled to generate large numbers of hypothetical events ([LeVeque et al., 2016](#)); accordingly, the approach we use is based on a Karhunen-Loève (K-L) expansion to generate slip patterns with correlation lengths that are thought to be “reasonable” from studies of past events, e.g., ([Mai and Beroza, 2002](#); [Goda et al., 2016](#); [Melgar and Hayes, 2019](#); [Crempien et al., 2020](#)), as discussed further in Section 2. We stress, however, that we do not claim we have the “correct” distribution, or even the best possible based on available science, and so our focus is on a methodology that could also easily be applied to other choices of the PDF. We also suggest that any PTHA study intended as guidance for decision makers should include a sensitivity study that considers how robust the results are to changes in the assumed earthquake distribution, along with other approaches for assessing the effects of the epistemic uncertainty inherent in this problem.

In this paper we focus on how best to handle the aleatoric uncertainty, i.e., assuming that we have a probability distribution to use for the PTHA, how do we efficiently create hazard curves based on this distribution? The brute force approach would be to choose a very large number N of samples from the distribution, perform a numerical tsunami simulation with each, and then (for each location of interest and each exceedance value) determine the number \hat{N} of samples for which h_{\max} exceeds \hat{h} . Then the ratio \hat{N}/N is an estimate of a conditional probability that h_{\max} exceeds \hat{h} given that some event from the set of all possible events occurs. If P_{total} is the annual probability that any event from the classes considered occurs, then $P_{\text{total}}\hat{N}/N$ could be used as the annual probability of h_{\max} exceeding \hat{h} .

The primary difficulty we address is that N may need to be very large in order to get meaningful statistics, particular for the relatively unlikely but most dangerous higher values of \hat{h} . Since a single tsunami simulation with a reasonable spatial resolution can take several minutes if not hours of computer time, this is problematic; and even more so if one also wants to do sensitivity studies and/or must consider many different communities over hundreds of miles of coastline.

A fundamental problem already arises when we ask for a reasonable value of P_{total} since it depends very much on what set of possible earthquakes to consider. Since earthquakes with magnitude less than Mw 7.5 rarely cause damaging tsunamis we could define P_{total} as the annual probability of any event with magnitude greater than this. In Section 2 we discuss our choice of PDF for the distribution of earthquakes. Although not necessarily correct for large subduction zone events, the Gutenberg-Richter law is a reasonable starting point for choosing a distribution. According to this, a magnitude Mw 7.5 event is 10 times more likely than an 8.5 event and roughly 32 times more likely than Mw 9.0. Hence, for example, if we sample $N = 3200$ events we would expect perhaps 100 to be Mw 9.0 or larger, a rather sparse sampling of these important and potentially quite diverse events. Moreover, most of the samples would be small events for which there is little or no inundation, a clear waste of computing resources.

We can address this by a simple application of *importance sampling*. We first split the space of all possible events of interest into a small number of *classes*. For illustration in this paper we use four classes based on magnitude: Mw 7.5, 8.0, 8.5, and 9.0, but this could easily be expanded. We then assign an annual probability to each class, call this P_j for class j ($j = 1, 2, 3, 4$ for our case, for which we assign the probabilities shown in Table 1). We then take N_j samples from class j , compute the fraction \hat{N}_j that exceed \hat{h} , and use $P_j\hat{N}_j/N_j$ as an estimate of the annual probability of exceeding \hat{h} by an event from class j . These can then be combined to obtain the annual probability of exceeding \hat{h} by any event (see Section 6). The advantage of splitting into classes is that we can choose a large number of events in a class corresponding to high impact but low probability (Mw 9.0 in our case) and then the corresponding fraction \hat{N}_j/N_j is weighted by a smaller annual probability P_j when combining with the probabilities obtained from other classes. For illustration, we have chosen to take $N_j = 500$ for each of the four classes so that we only consider 2000 events in total but 500 of them are in the Mw 9.0 class. (In Section 6 we discuss the rationale and implications of choosing this number of realizations.)

Next we tackle the problem that even 2000 tsunami simulations may be excessively demanding, particularly when we expect that many of these events will give very similar inundation patterns and depths as other events, and so in principle it should be possible to estimate the hazard curve with fewer simulations of judiciously chosen representative tsunamis. We develop an approach for *clustering* that can be applied to the 2000 events before doing any fine-scale tsunami modeling, in order to identify clusters of events that we expect to give very similar tsunami impact in the location of interest. We then do a fine-scale tsunami model of only one realization from each cluster (which we call the “cluster representative”) and assign it a weight that is based on the collection of events in that cluster. Based on this we can estimate the

contribution that this cluster should make to each hazard curve. This clustering is explained in much more detail in Section 5.

The clustering approach we illustrate in this paper is based on doing a coarse-grid tsunami simulation for each of the 2000 realizations, with a grid resolution that allows much faster simulation, but is too coarse to properly represent the tsunami inundation over the communities of interest. However, we show that these coarse grid simulations give information in the form of proxy variables that can be used to very effectively cluster the events.

Moreover, the coarse grid simulations can be greatly enhanced to provide “pseudo-fine” results that are at the resolution of the desired fine grid and that agree very well with the actual fine-grid simulations of the same realization, but are much cheaper to compute. This enhancement is performed in part using information about the difference between the coarse and fine grid simulations performed for the few realizations where both are available (the cluster representatives). This procedure is described in more detail in Section 4.

For illustration we consider two sample communities: Crescent City, CA, which is near the southern extent of the CSZ and Westport, WA, roughly 570 km north of Crescent City. These communities are both at high risk to CSZ tsunamis and have been the subject of past studies. They also have quite different topographic features as discussed further in Section 3. The same set of 2000 CSZ realizations was used for each site, although the clustering algorithm is applied separately to each, since the set of realizations that give similar inundation patterns at one site may not form a suitable cluster at the second site. For illustration we show that selecting only 18 clusters (and hence performing only 18 fine-grid simulations for each site) gives hazard curves and maps that compare very well with those obtained if all 2000 realizations are simulated on the fine grid, particularly after adding in additional information obtained from coarse-grid simulations of each realization. These results are presented in Section 7.

Some of the techniques presented in this paper were first developed as part of a project funded by FEMA Region IX, and presented in the project Final Report by Adams et al. (2017). Subsequently we have improved some of these techniques. We are also now using a probability distribution that is potentially more realistic than the original choice, and we consider two different target communities with quite different topography in order to better test the general applicability of these ideas. The original report and associated webpages (Adams et al., 2017) contain more discussion of some of these ideas, along with illustrations of some related approaches that are not reported in this paper.

Research on PTHA using stochastic collections of sources goes back many years, see for example the early review Geist and Parsons (2006) and the more recent ones of Geist and Lynett (2014) and Grezio et al. (2017) for many more references.

Recently, several researchers have adopted the use of a K-L expansion to generate large suites of realizations for PTHA studies of particular regions and/or to study sensitivities and uncertainty. For example, Gonzalez et al. (2020) generated 400 realizations for a hybrid deterministic/probabilistic tsunami hazard assessment to Iquique, Chile, and Crempien et al. (2020) generated 10,000 realizations on a idealized fault and performed GeoClaw tsunami simulations of each on idealized topography to study the effect of spatial slip correlation on tsunami intensity. The techniques developed in this paper could help to accelerate such studies.

Research on reducing the work required to handle large sets of realizations has also been done by others. For example, Sepulveda et al. (2017) used the K-L expansion together with a stochastic reduced order model to obtain better results than with a brute force Monte Carlo simulation. Even closer to the

methodology presented in this paper is the source filtering approach developed recently by Volpe et al. (2019). They started with a suite of more than 1.7 million scenarios that affect their study region, and used a clustering algorithm based on cheaply obtained proxies to reduce this to a smaller set of 1154. They then performed fine-grid simulations for one representative from each cluster. One significant difference in our approach is that we use coarse-grid simulations (using the full nonlinear tsunami model, including onshore inundation over the study region) and the associated pseudo-fine results, which allowed us to obtain PTHA results with fewer fine-grid simulations. On the other hand, we performed 2000 coarse-grid simulations to obtain these, whereas Volpe et al. (2019) performed the clustering based on proxy data that was more cheaply obtained. A hybrid approach might be to apply our methodology to the 1154 cluster representatives identified in (Volpe et al., 2019), performing only coarse-grid inundation simulations of these, and then further clustering into a much smaller set for the fine-grid simulations.

2 EARTHQUAKE PROBABILITY DENSITY AND REALIZATIONS

Probability distributions proposed for Cascadia Subduction Zone (CSZ) earthquake magnitudes have included both characteristic and Gutenberg–Richter (G-R) types. More generally, Parsons et al. (2012) noted that this is a long-standing controversy for many other fault zones. Consequently, they developed both types of distribution models for the Nankai Trough, based on data from many past events. The characteristic earthquake model was based on fixed rupture geometries and historical/palaeoseismic recurrence times, and the G-R model was based on fault-slip rates and an estimated distribution slope (b-value). They found that the G-R distribution replicated much of the spatial distribution of observed segmented rupture rates along the Nankai, Tonankai and Tokai subduction zones, although with some rate differences between the two methods in the Tokai zone. Thus, where supporting information exists (e.g. palaeoseismic and historical recurrence data), and fault segmentation observations are absent, they suggested that very simple earthquake rupture simulations based on empirical data and fundamental earthquake laws could be useful forecast tools in settings with sparse data from past events. We thus view a G-R distribution of magnitudes as adequate for this study.

We generated 2000 slip realizations over 4 magnitude classes: M_w 7.5, 8.0, 8.5, and 9.0 (with 500 of each). To determine the annual probabilities of earthquakes in each of the magnitude classes, we follow a G-R law using a b-value of 1, indicating "normal" seismic behavior. We also assume a yearly rate of occurrence of a M_w 9.0 along the CSZ as once every 526 years based on paleotsunami records from Goldfinger et al. (2012). This implies an a-value of 6.279 in the G-R relation and gives us annual probabilities P_j in Table 1 for each of our magnitude classes. Incidentally, Table 1 could also be extrapolated to show that the CSZ should have a $M_{6.3}$ earthquake every year. However, this is not the case on the anomalously-quiet CSZ. Nonetheless, for the purposes of presenting a PTHA methodology, using a Gutenberg-Richter law is a reasonable starting point.

Class j	M_w	Annual probability P_j	Return period (years)
1	7.5	0.06012	16.6
2	8.0	0.01901	52.6
3	8.5	0.00601	166.3
4	9.0	0.00190	526.0

Table 1. Table of the four magnitude classes used in this work, with the annual probability P_j for an event from each class, along with the corresponding return period $1/P_j$, based on the Gutenberg-Richter formula $p = 10^{6.279 - M_w}$ for a magnitude M_w event.

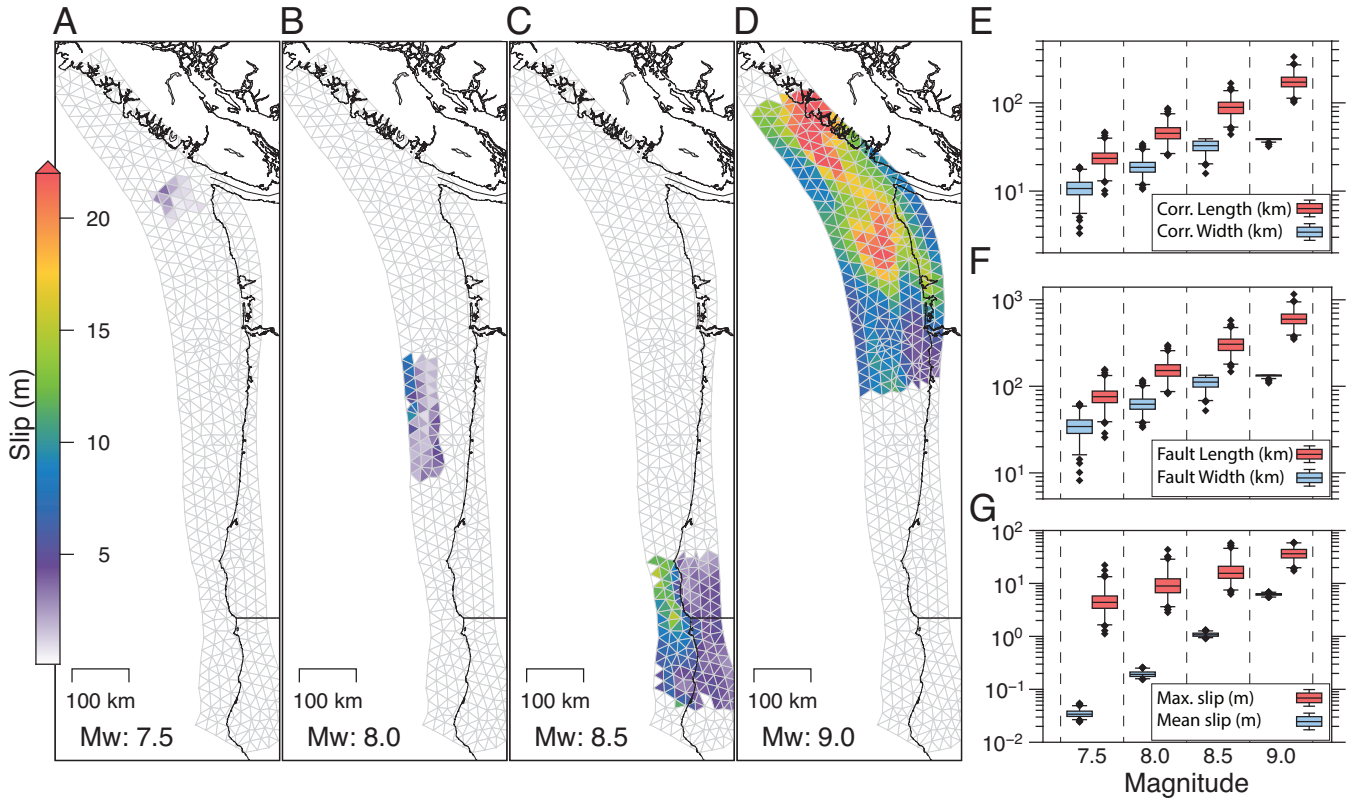


Figure 1. Panels A-D: Example earthquake realizations for each magnitude class. E: Variability in correlation length and width per each magnitude. F: Variability in total fault length and width per magnitude class. G: Variability in mean and maximum slip, in meters, for realizations in each magnitude class

While a full PTHA analysis should include a wide range of potential sources, here we limit our earthquake realizations to imitate a series of thrust events located on the megathrust interface along the CSZ. To introduce variability to each realization, we allow for geophysically reasonable variations in slip distribution, location, and rupture dimension. An example of a rupture from each magnitude class is shown in Figure 1. We employ a regional fault geometry that approximates the CSZ from McCrory et al. (2012). This is then discretized into triangular subfaults using the three-dimensional finite element solver GMSH (Geuzaine and Remacle, 2009). A triangular mesh allows a variable strike and dip that can better approximate the McCrory et al. (2012) geometry than a rectangular discretization. Our area of interest extends along the entire CSZ margin and down to a depth of 30 km beyond which slip is not expected to continue (Frankel et al., 2015).

The fault length, L , and width, W , of each realization is first defined using a magnitude dependent dimension scaling law that is specific to subduction zones from Blaser et al. (2010) where:

$$\begin{aligned}\log_{10} L &= -2.37 + 0.57M_w \\ \log_{10} W &= -1.86 + 0.46M_w.\end{aligned}\tag{1}$$

In order to introduce variability in ruptures of the same magnitude, as is observed from past earthquakes, the length and width of each realization is slightly augmented. We sample from a normal probability density function centered about the length and width parameters from above using the standard deviations provided in Table 2 of Blaser et al. (2010). The rupture extent is then centered about a randomly chosen subfault

within our CSZ mesh geometry. If the chosen subfault is located in such a place that the rupture extent exceeds the bounds of the rupture geometry, then it is moved up/down dip and/or along strike until it falls completely within the CSZ.

Once the bounds of the rupture area are established, we generate a stochastic slip distribution using an application of the Karhunen-Loève (K-L) expansion following [LeVeque et al. \(2016\)](#) and [Melgar et al. \(2016\)](#). Here, we assign slip over participating subfaults using a von Karman correlation function, $C(r)$, which replicates the statistics of slip distributions as observed from finite-fault solutions of past moderate sized earthquakes ([Mai and Beroza, 2002](#)). Here, the correlation between the s th and d th subfaults (in along strike and along dip directions) is

$$C_{sd}(r_{sd}) = \frac{G_H(r_{sd})}{G_0(r_{sd})} \quad (2)$$

where

$$G_H(r_{sd}) = r_{sd}^H K_H(r_{sd}). \quad (3)$$

H is the Hurst exponent (set in this study as 0.75), K_H is the modified Bessel function of the second kind, and r_{sd} is a length measurement for s th and d th subfaults that depends on the distance between subfaults in the along strike (r_s) and along dip (r_d) directions as well as the correlation length along strike (a_s) and dip (a_d), written as

$$r_{sd} = \sqrt{\frac{r_s^2}{a_s} + \frac{r_d^2}{a_d}}. \quad (4)$$

The correlation length and width for each realization governs the size of asperities and uses a magnitude dependent scaling law from [Melgar and Hayes \(2019\)](#) where

$$\begin{aligned} a_s &= 17.7 + 0.35L_{\text{eff}} \\ a_d &= 6.7 + 0.41W_{\text{eff}} \end{aligned} \quad (5)$$

The variables L_{eff} and W_{eff} are based on the effective length and width scales (in kilometers) from equation (1) and are determined from [Mai and Beroza \(2000\)](#). Using the defined correlation function the distribution of slip across the fault surface is treated as a spatially random field. The vector \mathbf{s} containing the amount of slip at each subfault is then given by the Karhunen-Loève expansion as,

$$\mathbf{s} = \mu + \sum_{k=1}^{n_f} \mathbf{z}_k \sqrt{\lambda_k} v_k, \quad (6)$$

where μ is the mean slip over the entire fault, λ_k and \mathbf{v} are the eigenvalues and eigenvectors of the chosen correlation function. z_k are random numbers normally distributed with a mean of 0 and standard deviation of 1. There are as many eigenvectors or eigenmodes (n_f) as there are subfaults in the model geometry. Equation (6) is a statement of how each eigenmode distributes slip away from the background mean model modulated by a random number z_k and thus achieving a stochastic realization of slip. For tsunami modeling because seafloor deformation is a relatively long period phenomena [LeVeque et al. \(2016\)](#) showed how high order modes contribute relatively little to tsunamigenesis so it is possible to truncate the sum to just a

few tens or even hundreds of terms. Here we limit the number of contributing modes to no more than 200. Finally for the background mean model we assume enough homogeneously distributed slip to match the target magnitude given the chosen fault dimensions. It is also possible to make other choices for μ such as a known slip distribution or a geodetic locking model (Goldberg and Melgar, 2020).

We cap the upper level of slip possible for any realization in this study to 60 m, as was recommended in Melgar et al. (2016) and based on plate convergence rates from McCaffrey et al. (2007). This choice allows for the possibility that the most recent large earthquake, the 1700 M9 Cascadia event, may not have released all strain on the megathrust interface. This cap is in place in order to limit the possibility of unrealistically large amounts of slip in any earthquake realization. It should be noted that this cap in fault slip creates an upper limit in tsunami intensity that may be reflected in any final PTHA analysis.

We calculate the total seafloor deformation of each earthquake realization using angular dislocations for triangular subfaults in an elastic half space (Comninou and Dundurs, 1975). This method is an adaptation from the Okada equations, which focus on rectangular subfaults (Okada, 1985). We obtain the deformation over the entire CSZ study region with at a 30" spatial resolution. This is a fine enough spacing to ensure that we recover slip features that may be present at our smallest magnitude class, including rupturing asperities at the smallest reported slip correlation length (see examples in Figure 2).

Seafloor deformation is directly translated to a disturbance at sea level by assuming an incompressible water column. While some large magnitude earthquakes can have rupture durations extending multiple minutes, this kinematic effect on the tsunami in the near-field is minimal (Satake, 1987; Williamson et al., 2019). Here, we simplify the rupture process by treating all seafloor deformation as instantaneous and occurring at the initial time step of our tsunami model. It is this initial disturbance that initializes the tsunami model, as discussed further in Section 3.

Figure 2 shows both slip on the fault and the resulting seafloor deformation for two of the magnitude 9.0 realizations. We use R_i to denote the i th realization and the figure shows the realizations R_{1655} and R_{1999} out of the $N = 2000$ realizations, chosen because R_{1655} has slip concentrated on the southern margin of the fault while R_{1999} has it concentrated to the north, and hence they have very different effects in Westport and Crescent City, as discussed in the next section. This illustrates that even within a single magnitude class there are significant variations between the tsunamis generated.

3 COMMUNITIES OF INTEREST AND TSUNAMI MODELING

We consider two sample communities as shown in Figure 3. Crescent City, CA was used in our previous work on this topic (Adams et al., 2017), and was the subject of a previous PTHA analysis by González et al. (2014) and Adams et al. (2015). Tsunamis tend to focus in Crescent City due to the offshore bathymetry and harbor (Horrillo et al., 2008), and the central business district is bounded by the harbor, the low-lying Elk River valley to the east, and higher hills to the north.

Westport, WA lies on a peninsula at the entrance to Grays Harbor. The topography is below roughly 10m everywhere on the peninsula, and a number of north-south running ridges protect some areas from the direct waves arriving from the west that may still be flooded from the east after the tsunami enters Grays Harbor. Westport is the site of the Ocosta Elementary School, recently rebuilt to include the first tsunami vertical evacuation structure constructed in the US, for which tsunami modeling was presented by González et al. (2013).

Tsunami simulations are performed using GeoClaw Version 5.7.0, distributed as part of the open source Clawpack software (Clawpack Development Team, 2020). This solves the two-dimensional depth-averaged

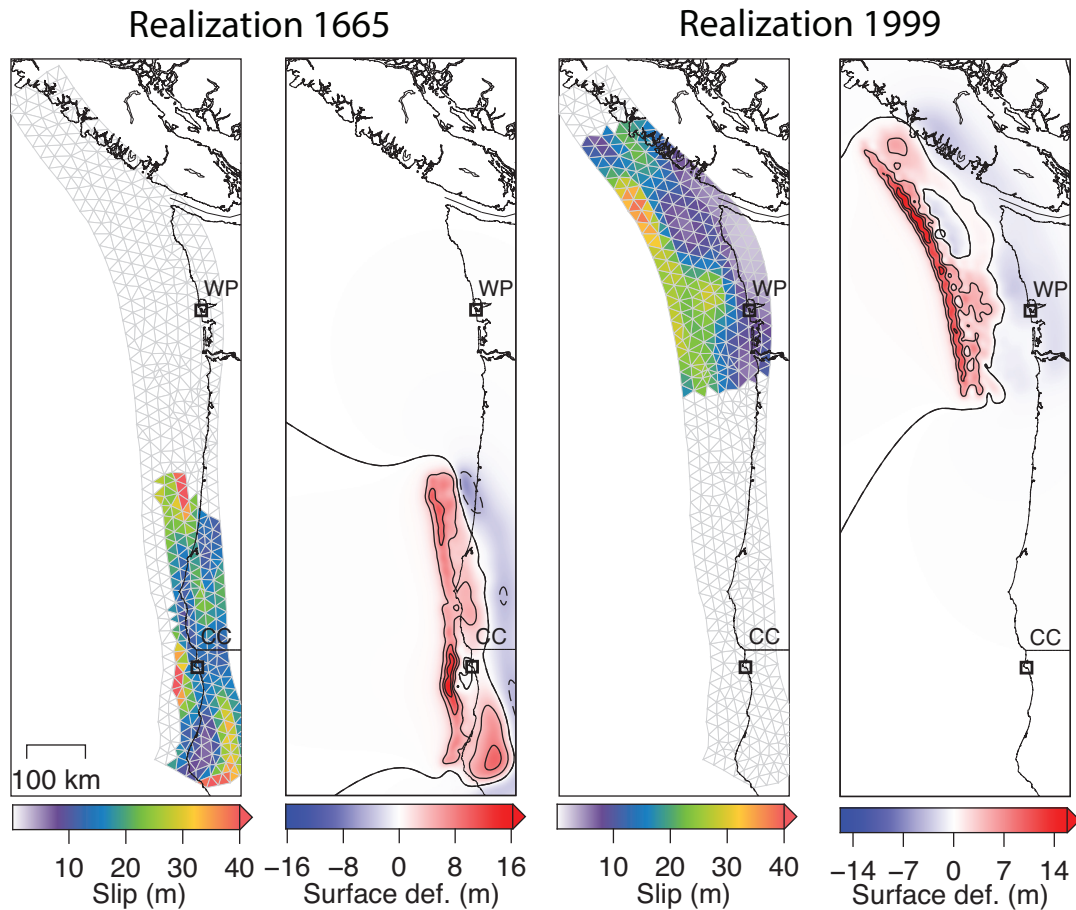


Figure 2. Slip on our fault geometry and associated seafloor deformation for two sample realizations, numbered 1665 and 1999, shown in Figures 4, 5, 6, and 7. WP marks the location of Westport, WA. CC marks the location of Crescent City, CA.

non-linear shallow water equations using adaptive mesh refinement on rectangular grid patches (in longitude-latitude coordinates). GeoClaw allows each cell to be wet or dry and to change dynamically, so that the wet/dry boundary of the coastline evolves as the tsunami inundates the coastal site of interest.

For this study we simulated the tsunami from each of the 2000 realizations in two separate simulations. The first set were the “fine-grid runs” where refinement down to 1 arcsecond (roughly 30 m in latitude, less in longitude) was enforced over both study sites. This provided the “ground truth” hazard curves and maps to use for comparison purposes, i.e., we assume that our goal is to produce good approximations to these curves and maps with much less work than was required to run 2000 fine-grid simulations. The second set of simulations were the “coarse-grid runs” in which the refinement only went down to 9”, and hence a factor of 81 fewer grid points on the finest level than in the fine-grid runs. Moreover on these coarser grids it is also possible to take larger time steps (while still respecting the Courant-Friedrichs-Lewy (CFL) condition required by the explicit finite volume method used in GeoClaw), potentially giving another factor of 9. However, since some of the computation takes place on coarser grids over the entire computational domain, the coarse grid simulations are typically only 5–20 times faster than the fine grid simulations; see below. We also note that for a real PTHA we might want to use even finer grids, e.g. 1/3” is often used now used for hazard studies, and 1/9” topography is becoming available in many locations. In this case the relative speedup for coarse-grid simulations could be much more dramatic.

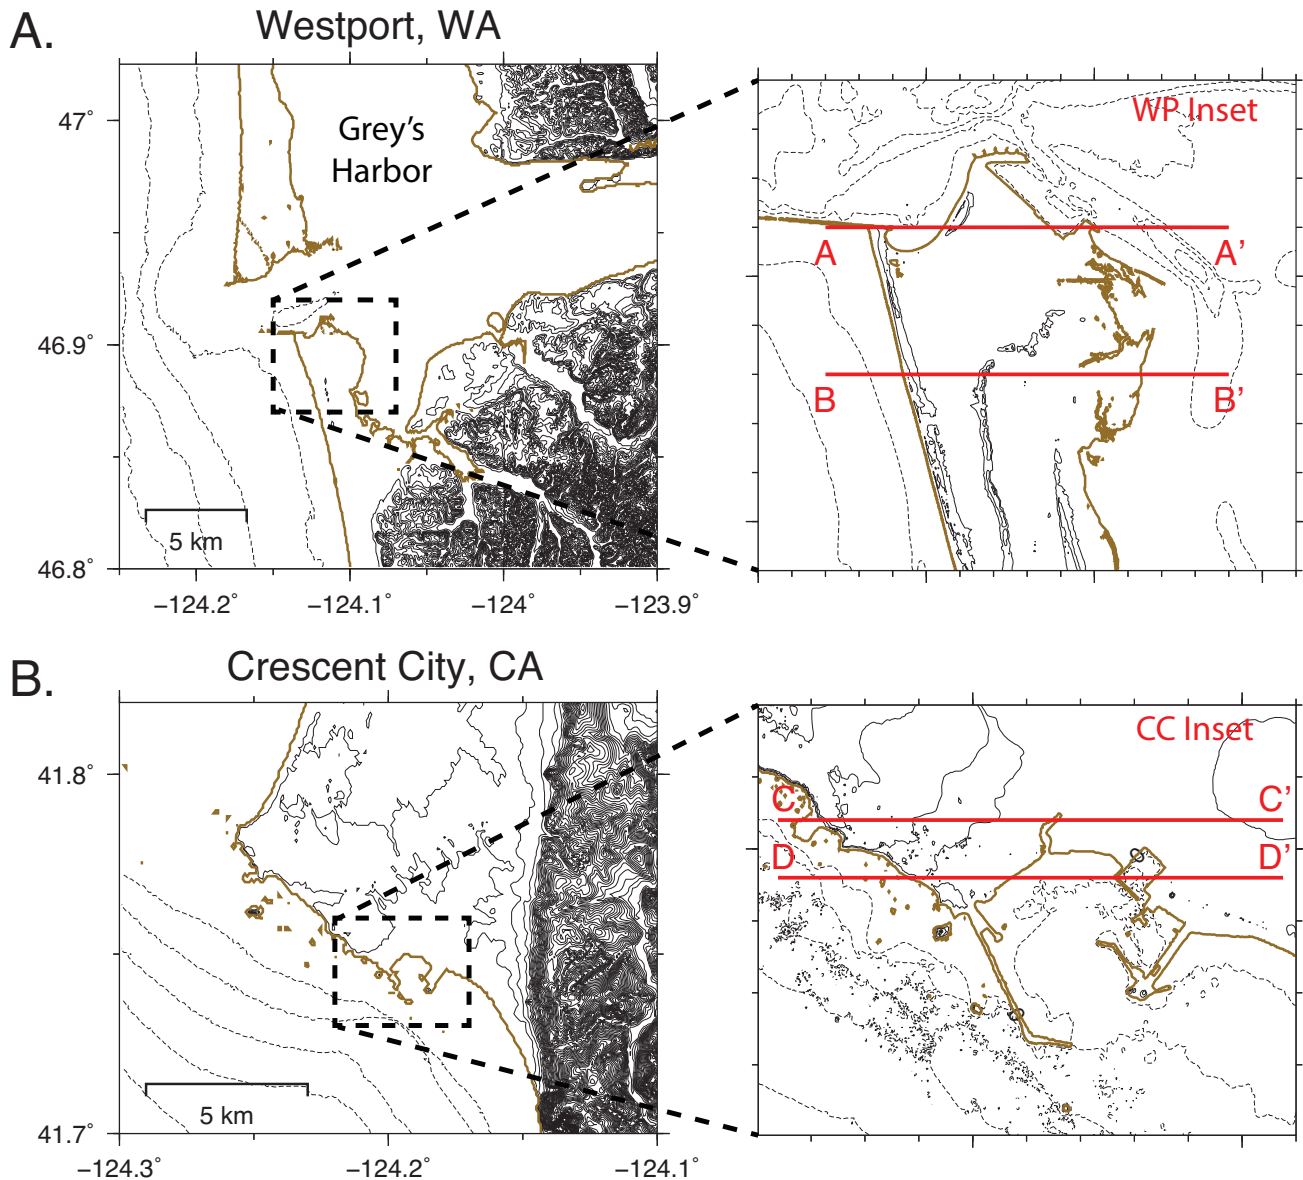


Figure 3. Communities of interest in this study A. Regional and inset view of Westport, WA and two representative cross sections shown in Section 7. B. Regional and inset view of Crescent City, CA with two representative cross sections. Both regional views have a bathymetric (dashed) and topographic (solid) contour interval of 10m. Inset figures use a contour interval of 5m. The coastline is differentiated with a bolded brown line.

We use adaptive mesh refinement to optimize the computational cost of each tsunami simulation. All simulations used three levels of refinement in the open-ocean, with grid resolution 1° , $6'$, and $3'$, and with regridding every few time steps to follow the propagating waves (based on a tolerance on the sea surface elevation). On the continental shelf, refinement is allowed to the next AMR level at $90''$. An additional refinement level of $9''$ is enforced around both study sites. For the coarse-grid simulations only these 5 levels of AMR are used. For the fine-grid runs, two additional AMR levels are introduced at $3''$ and $1''$ resolutions, and the study areas are forced to be resolved at the finest $1''$ level. The ETOPO1 topography DEM at 1 arcminute resolution (Amante and Eakins, 2009) was used over the full computational domain. A subset of the Astoria, Oregon $1/3''$ DEM (NOAA NCEI, Accessed 2017) was used around Westport, and around Crescent City a version of the Crescent City, California $1/3''$ DEM (NOAA NCEI, Accessed 2012)

was used that was modified to remove the pier in the harbor, since water flows under the pier, for an earlier PTHA study of this region by [González et al. \(2014\)](#).

In each simulation we monitor the maximum water depth h over a grid of points covering the study area (at the finest resolution of the simulation) over the duration of the simulation. For this study we ran each tsunami simulation to 4 hours of simulated time after the instantaneous seafloor displacement. Examining the results we found that in a few cases there were still significant edge waves trapped along the coast that could have lead to slightly larger values of the maximum at some points, so a realistic PTHA should run some realizations out to later times. For the purposes of this study our reference solution uses the maximum h over the same time period as our approximations and so comparisons are still valid.

At each point where h is monitored, these maximum values (denoted by h_{\max}^f for the fine-grid runs) are used to compute a reference hazard curve. The coarse-grid simulations produce their own set of h_{\max}^c values on a coarser set of points (which can be extended to the fine grid by piecewise constant interpolation within each coarse grid cell). These coarse-grid values are used both in the clustering algorithm and in computing a pseudo-fine result from each coarse result, as explained in the sections below.

The tsunami simulations were performed using the OpenMP feature of GeoClaw using 30 threads on a Linux server. The total CPU time varied for each realization, depending on whether the initial deformation came from a small localized slip patch (requiring small regions of refinement in the ocean and possibly resulting in a negligible tsunami) or a larger rupture. Total CPU time (summed over all threads and over all 2000 simulations) was 49.2 hours for the coarse-grid runs and 255.6 hours for the fine-grid runs.

3.1 Sample realizations

Before proceeding, we first show h_{\max} in each study region for the two sample $M_w = 9.0$ realizations shown in Figure 2. Figures 4 and 5 show h_{\max} for R_{1665} at Westport and Crescent City, respectively, and similarly Figures 6 and 7 show h_{\max} for R_{1999} . In each case, panel A shows the fine-grid result h_{\max}^f while B shows the coarse-grid result h_{\max}^c . Note that the 9" grid cell resolution is clearly visible in B and that this coarse grid cannot resolve all features of the flow, but that the general order of magnitude is correct. Panel C shows the difference between coarse and fine results, which are substantial in some regions.

The remaining panels of each figure show the result of enhancing the coarse-grid results using techniques developed in the next section, where these will be discussed in more detail.

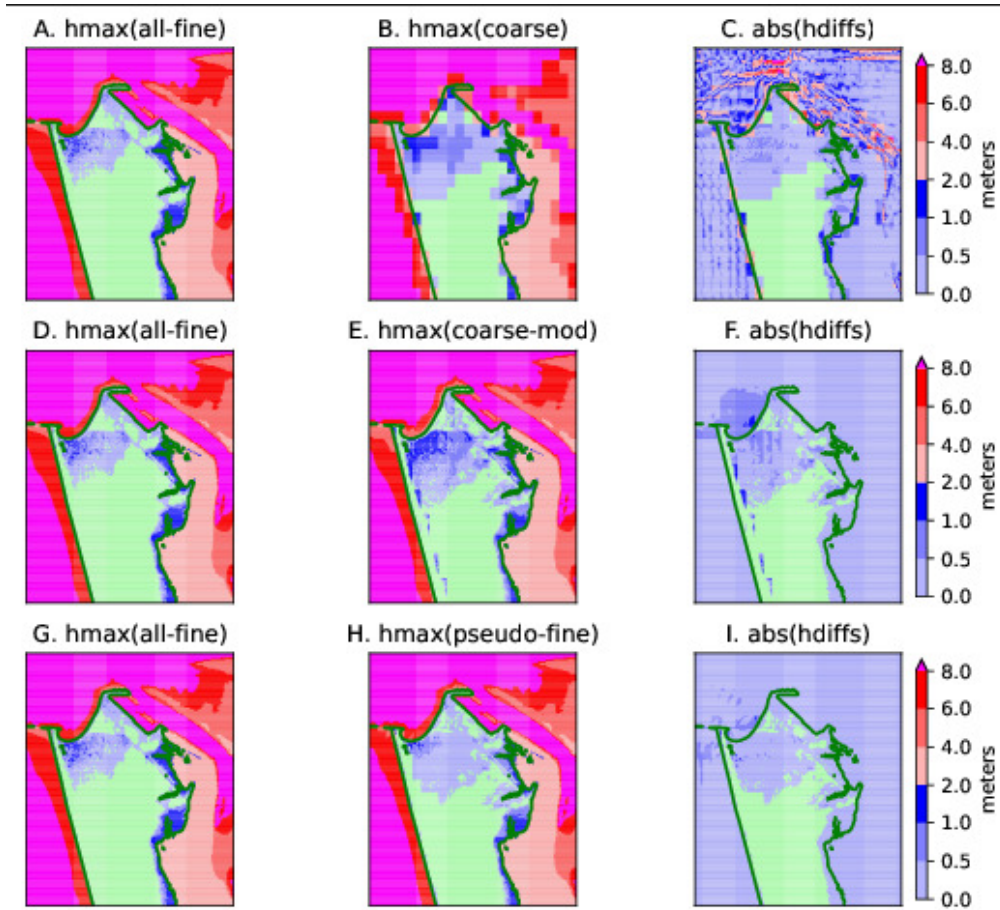


Figure 4. Sample results for realization 1665 at Westport, where the tsunami was small. (A,D,G) h_{\max}^f . (B) h_{\max}^c . (E) h_{\max}^c after coarse-mod corrections. (H) h_{\max}^c after pseudo-fine corrections. (C,F,I) Errors relative to h_{\max}^f . See the text for more explanation.

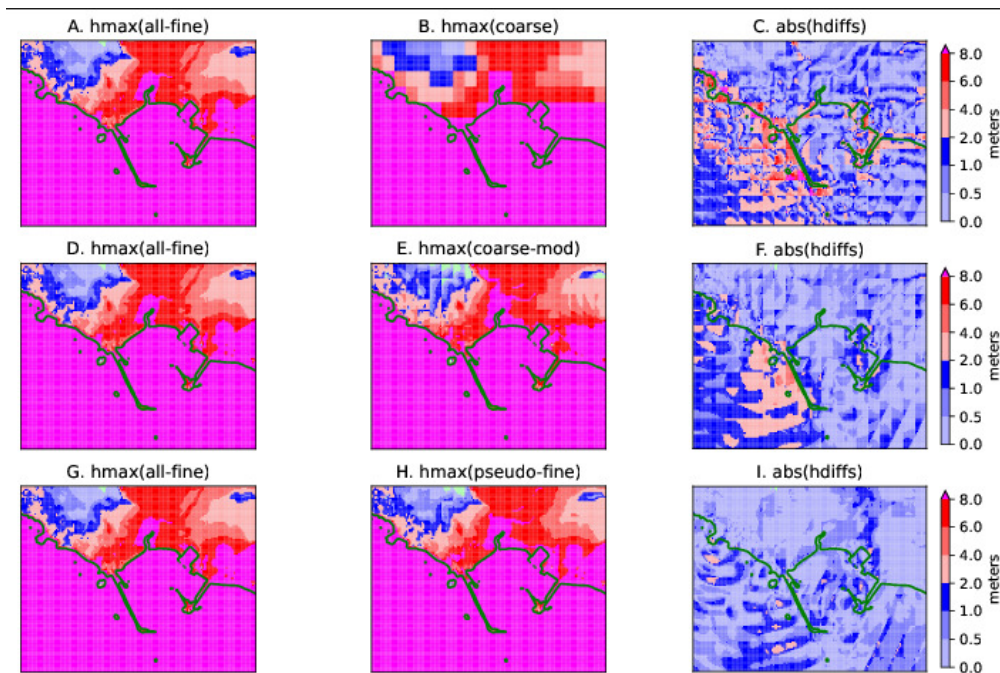


Figure 5. Sample results for realization 1665 at Crescent City, where the tsunami was large. Panels as described in Figure 4.

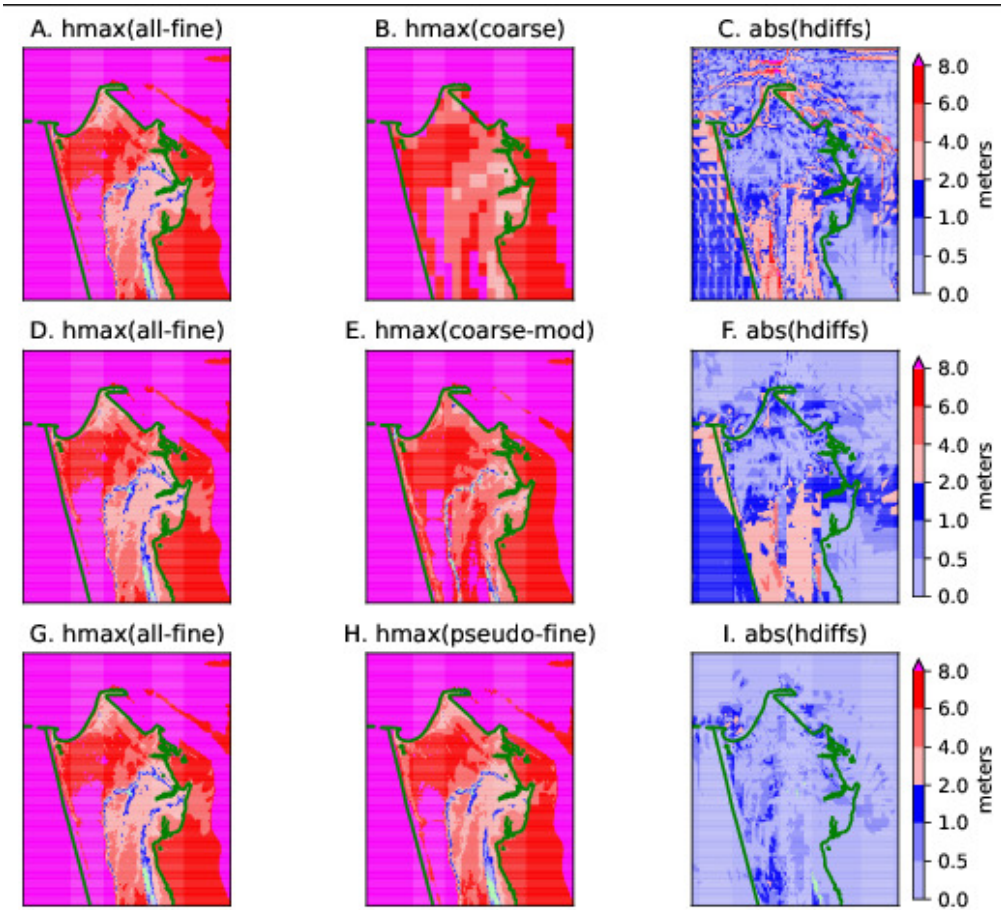


Figure 6. Sample results for realization 1999 at Westport, where the tsunami was large. Panels as described in Figure 4.

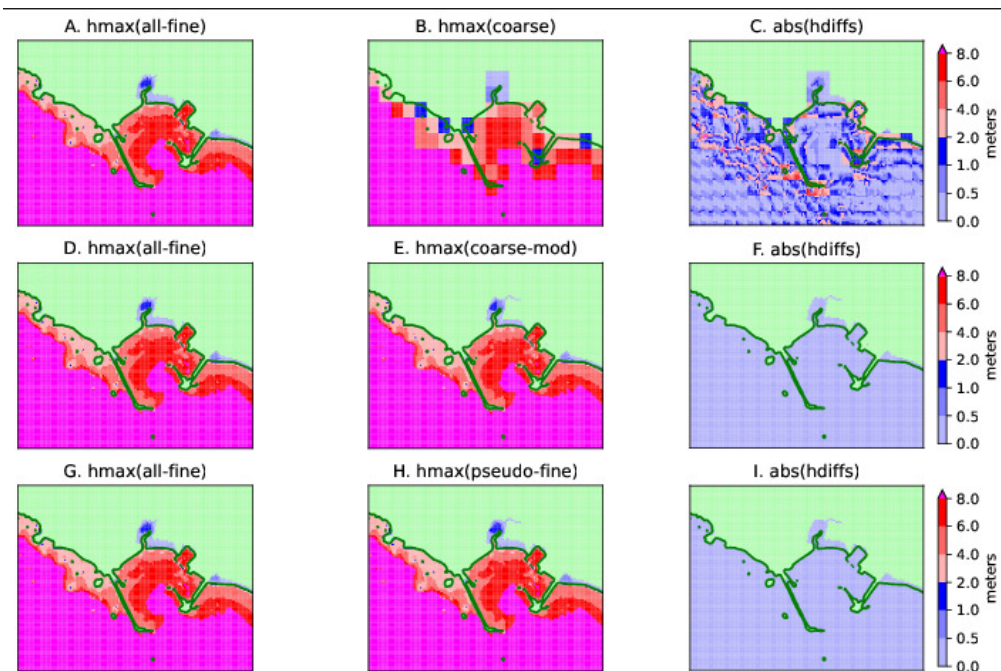


Figure 7. Sample results for realization 1999 at Crescent City, where the tsunami was small. Panels as described in Figure 4.

4 COARSE-MOD AND PSEUDO-FINE ENHANCEMENTS

Our PTHA approach starts by sampling N realizations, which we denote by R_i (for $i = 1, 2, \dots, N$). These may consist of N_j realizations from class j as described in Section 1, with $N = \sum_j N_j$. Performing coarse-grid simulations of each gives us h_{\max}^c at each location on a coarse grid covering the study region. We wish to avoid doing fine-grid simulations of all realizations, and instead we will use a clustering approach, described in detail in Section 5 below, to group these into K clusters and to choose one representative realization from each cluster. This “cluster representative” will be denoted by \bar{R}_k for the particular realization from cluster $k = 1, 2, \dots, K$.

One approach to approximating the hazard curves for N realizations using K clusters is to perform only one fine-grid simulation for each cluster (for a “cluster representative” realization selected from the cluster), and assign a weight to each that is the sum of the weights of all realizations in the cluster. We show results of this approach in Section 7. However, using only K events will give a hazard curve with only K jump discontinuities and cannot well approximate the true hazard curve if K is much smaller than N , particularly at the lower probabilities. In our example application, $N = 2000$ and we will choose $K = 18$.

Much better results are obtained if we also make use of the remaining $N - K = 1982$ coarse-grid simulations that were performed to do the clustering. The coarse grid results alone do not give sufficient resolution of h_{\max} for use directly, but they can be enhanced to approximate the inundation that each would produce on a fine grid with much less work than required to do the fine-grid simulation. This is done as a two-stage process. In the first (“coarse-mod”) step the coarse grid results are combined with the fine-grid topography to give estimates of the maximum depth on the fine topography. This is independent of the clustering and can be done immediately following each coarse grid simulation. The second (“pseudo-fine”) step uses the clustering, and the idea that the difference between the coarse-mod and fine-grid simulations at the cluster representative (both of which are available) gives a good indication of how other coarse-mod results in the same cluster should be adjusted to better approximate the result of a fine-grid simulation. We describe each of these in turn.

4.1 Modified coarse grid corrections

Each of the $N = 2000$ coarse-grid runs provides an estimate of h_{\max}^c at a set of coarse grid points covering the study region. These were computed using a coarse grid (with resolution 9” in our case), which is inadequate to represent the community of interest. For example, the top row of Figures 4–7 shows h_{\max} in each study region for two particular realizations, and the difference in resolution is apparent between the fine-grid simulation (h_{\max}^f in Panel A of each figure) and the corresponding coarse-grid simulation (h_{\max}^c in Panel B). Panel C of each figure shows the difference between the fine and coarse results, after the coarse-grid results are interpolated to the fine grid as a piecewise constant function over each coarse grid cell.

The coarse-mod corrections are based on the observation that the maximum water surface $\eta_{\max} = B + h_{\max}$ (where B is the pre-seismic topography) is often much more smoothly varying over a community than is the maximum water depth h_{\max} . (Note that even a wave filling in to a constant η_{\max} throughout the study region would still have large variations in $h_{\max} = \eta_{\max} - B$ due to the variations in topography B .) Hence at any point, if we assume η_{\max} is roughly correct, we can get a better estimate of h_{\max} by subtracting off the fine-grid topography at this point from the η_{\max}^c value predicted by the coarse-grid simulation. As usual, we focus on values at a single grid point on the fine grid. Let B^f represent topography from the fine-grid simulation at this point and B^c the topography value from the coarse-grid simulation in the coarse cell containing this point. Then $\eta_{\max}^c = B^c + h_{\max}^c$. The correction we make defines a modified

value h_{\max}^{cm} at this point as

$$h_{\max}^{cm} = \eta_{\max}^c - B^f = B^c + h_{\max}^c - B^f = h_{\max}^c - \Delta B \quad (7)$$

where

$$\Delta B = B^f - B^c. \quad (8)$$

In other words, we simply adjust h_{\max}^c at each fine grid point by ΔB , the difference between the fine and coarse topography at this point.

However, there are a few special cases where we can not use (7) above to set h_{\max}^{cm} . Clearly, if $\Delta B > h_{\max}^c$, we can not allow h_{\max}^{cm} to become negative (the water reaches the coarse bathymetry level, but not the fine level), so h_{\max}^{cm} is set to 0 in this case. Another special case is when $h_{\max}^c = 0$. In this case, water may or may not reach the fine bathymetry level. To determine if it does, we define an η threshold value, called η_T and now use η_c to denote the η_{\max} value at the point of interest on the coarse grid. Over the four neighboring grid cells around the coarse cell containing the location of interest, we find η_T , the maximum η^c value where $h_{\max}^c > 0$; that is, the threshold where we have seen flooding locally. Then we set $\eta^{cm} = \eta_T$ ($h_{\max}^{cm} = \eta_T - B^f$), $\eta^{cm} = \eta^c$ ($h_{\max}^{cm} = B^c - B^f$), or $\eta^{cm} = B^f$ ($h_{\max}^{cm} = 0$) depending on whether $B^f < \eta_T \leq B^c$, $\eta_T > B^c$, or $\eta_T \leq B^f < B^c$, respectively. With these definitions, we define the “coarse-mod” correction h_{\max}^{cm} to be $\eta^{cm} - B^f$. Based on the above definitions this gives:

$$h_{\max}^{cm} = \begin{cases} h_{\max}^c - \Delta B & \text{if } 0 \leq \Delta B < h_{\max}^c \\ 0 & \text{if } 0 < h_{\max}^c \leq \Delta B \\ h_{\max}^c - \Delta B & \text{if } \Delta B < 0 < h_{\max}^c \\ \eta_T - B^f & \text{if } h_{\max}^c = 0 \text{ and } B^f < \eta_T \leq B^c \\ B^c - B^f & \text{if } h_{\max}^c = 0 \text{ and } B^f < B^c < \eta_T \\ 0 & \text{if } h_{\max}^c = 0 \text{ and } \eta_T \leq B^f < B^c \end{cases} \quad (9)$$

The second row of Figures 4–7 shows examples of the effect of this. Panel D is again the fine grid h_{\max} but now Panel E shows the h_{\max} estimated on the fine grid after applying these coarse-mod corrections. Panel F shows the resulting errors.

An earlier version of this coarse-mod strategy was used in (Adams et al., 2017), and more examples are given in the Appendix of that report showing how well the modified coarse data can compare to the original coarse data and the fine data. We have since improved this strategy by looking only locally for the threshold value η_T as opposed to across the entire community in the FEMA report, and have allowed this value to even be negative, which makes the strategy more applicable to locations near the shoreline that see little inundation during an uplift event. These improvements also result in a more effective pseudo-fine strategy as discussed below.

4.2 Pseudo fine grid corrections

We now present an approach to better improving each of the coarse-mod results by (9) by also using the fine-grid runs performed for each cluster representative. We begin by clustering the N modified coarse grid runs (or the original N coarse grid runs) into a small number of non-overlapping clusters. For this paper, the clustering was done using the original $N = 2000$ coarse grid runs, to produce $K = 18$ clusters, as described in Section 5. These clusters contain different numbers of runs. Each cluster has one run designated as its cluster representative, which we will denote by \bar{R}_k for the particular realization from cluster $k = 1, 2, \dots, K$.

After the clustering has been done, more information is available that can be used to further improve the coarse-mod approximations. Since we assume that the cluster representative \bar{R}_k is somewhat typical of the pattern of flooding seen for all realizations in the cluster, and since we have both fine-grid and coarse-mod results available for this representative realization, we can use the difference between these as an estimate of what the difference between fine-grid and coarse-mod results would be for all realizations in the cluster. This is used to modify each coarse-mod result to get a better approximation to the expected fine-grid result. This is what we call the pseudo-fine result for each realization. We again use $h_{\max}^{cm}(R_i)$ to denote the h_{\max} value obtained for a particular realization R_i from the coarse-grid simulation after applying the coarse-mod corrections, and similarly $h_{\max}^f(R_i)$ comes from the fine-grid simulation, as always focusing on a single spatial location. Then the pseudo-fine approximation at this location for each realization R_i in Cluster k is given by

$$h_{\max}^{pf}(R_i) = h_{\max}^{cm}(R_i) + (h_{\max}^f(\bar{R}_k) - h_{\max}^{cm}(\bar{R}_k)). \quad (10)$$

Note in particular that the pseudo-fine result for the cluster representative itself (i.e., for $R_i = \bar{R}_k$) agrees exactly with the fine-grid result for that realization.

The third row of Figures 4–7 show examples of the effect of this. Panel G is again the fine grid h_{\max} but now Panel H shows the h_{\max} estimated on the fine grid after applying these pseudo-fine corrections. Panel I shows the resulting errors relative to the fine-grid result. Additional illustrations of this idea can be found in the Appendix to (Adams et al., 2017).

5 CLUSTERING

In this section we discuss how one can subdivide the individual events into a small number of clusters that are likely to have similar inundation patterns. The clustering will be based on proxy (surrogate) quantities for each event that can be computed solely from the coarse-grid (low-resolution) simulations, whose runtime is orders of magnitude smaller than the fine-grid (high-resolution) tsunami inundation simulations.

In our previous work (Adams et al., 2017), proxy variables based only on the seafloor deformation of each realization were also considered. These are much cheaper to compute than the coarse-grid simulation proxies, but did not do as good a job of clustering, even though in that work we only considered ruptures on the southern margin of CSZ and only one near-field study region, Crescent City. Given the wider range of events now being considered, where many events are localized far away from a study region, we believe that it may be harder to develop robust proxies based only on the seafloor deformations. However, due to the greater efficiency of that approach, this could be a fruitful area for future research. We also note that Volpe et al. (2019) used seafloor deformation near the study region, both to classify realizations into near-field and far-field, and also in clustering the near-field realizations.

In this work we use three proxy variables computed from the coarse-grid simulations for each realization. Each realization thus corresponds to a point in a three-dimensional space and various clustering methods can then be used to identify clusters. We consider variables that attempt to capture aspects of the spatial variation of the inundation patterns. For each realization, and each coastal location, we compute the value η^c from the coarse-grid tsunami simulation that measures the surface level at maximum inundation, defined by

$$\eta^c = B^c + h_{\max}^c, \quad (11)$$

where B^c represents the pre-event topography on the coarse grid. We consider the spatial variation of η over the onshore points that are flooded: the grid-cells in the simulation centered at (x_i, y_j) that satisfy

both $B^c(x_i, y_j) \geq 0$ and $h_{\max}^c(x_i, y_j) > 0$. We will denote these values from the coarse-grid simulation by η_{ij}^c . The total number of flooded onshore points will be denoted by N_{flood} .

As the proxy variables, we will use the following statistics of η , where the sums are over all onshore flooded points (i, j) :

$$\begin{aligned}\eta_{\text{logsum}}^c &= \log \left(1 + \sum_{ij} \eta_{ij}^c \right), \\ \eta_{\text{mean}}^c &= \frac{1}{N_{\text{flood}}} \sum_{ij} \eta_{ij}^c, \\ \eta_{\text{sd}}^c &= \sqrt{\frac{1}{N_{\text{flood}}} \sum_{ij} \left(\eta_{ij}^c - \eta_{\text{mean}}^c \right)^2}.\end{aligned}\tag{12}$$

The first variable η_{logsum}^c is a measure of the total extent and elevation of the flooding, while the second variable η_{mean}^c is the mean surface elevation. The third variable η_{sd}^c measures the spatial variation of surface elevation over the onshore flooded region. The first two variables summarize the severity of the flooding while the third variable summarizes the spatial variation of the flooding pattern.

Utilizing these quantities, the coarse-grid inundation from each individual realization can be mapped to a point in the three-dimensional space of proxy variables, which we will call “proxy-space”, as shown in Figure 8. To perform clustering it is also necessary to define a metric that measures the distance between two points in this space, and here we simply use the Euclidean distance (the square root of the sum of squares of differences in each of the three proxy variables). We then use K-means clustering (Lloyd (1982)), as implemented in `scipy.sklearn` by Pedregosa et al. (2011) to cluster the 2000 points in proxy-space into K clusters, with the property that each point belongs to the cluster with the closest centroid (as measured in the specified metric). Note that because η_{mean}^c typically has larger magnitudes than the other two proxy variables (see Figure 8), the use of the Euclidean distance effectively weights differences in η_{mean}^c more heavily than differences in the other proxy variables. We also tried first normalizing all proxy variables (equivalent to using a different weighted metric) and somewhat different clusters were generated but with very similar final results in the hazard curves. If using proxy variables with vastly different magnitudes, and/or where some variables are thought to be more important than others, some care should be used in choosing the metric.

This clustering is done independently for the two study regions, and the results are shown in the scatter plots in Figure 8. The plot also highlights which realization is closest to the centroid of each cluster, which we refer to as the cluster representative \bar{R}_k for Cluster k . The number of clusters is determined by the user, and in this instance 18 clusters were used to subdivide the 2000 events of varying magnitude. We also tried clustering with only 6 or 12 clusters but the results were not as good, while using more than 18 clusters did not seem necessary for this particular data set. Volpe et al. (2019) use an iterative procedure to select the number of clusters by enforcing a maximum distance between cluster members and the centroid and this might be a good approach in general.

From Figure 8 one observes that there is a general monotonic behavior with respect to all three variables. Higher magnitude events tend to have higher values for all three variables. But also note that the scatter plot for the two coastal communities show qualitatively different patterns. While the points in proxy space for Westport show a more predictable behavior with higher magnitude events more concentrated along a smooth, monotone increasing curve, there is much more variation in the proxy variable η_{sd}^c in Crescent

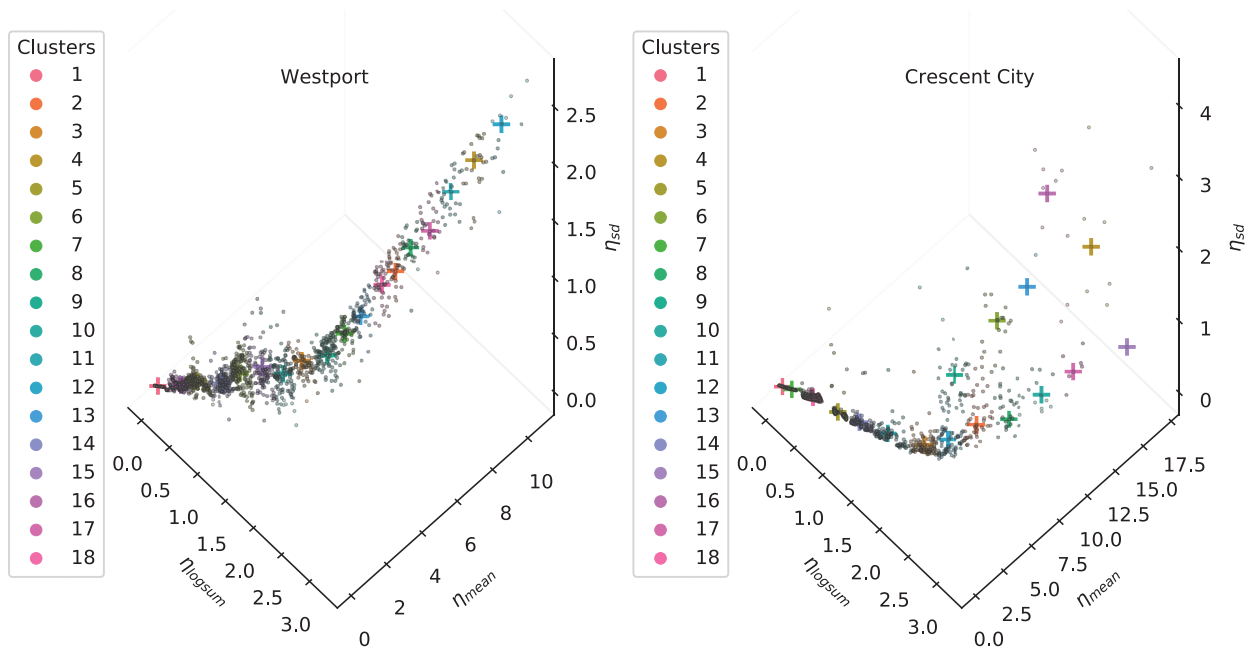


Figure 8. Clustering results for the two study regions. In each case all 2000 realizations are represented as a point in the three-dimensional proxy-space, and colored by cluster after with 18 clusters in each case. The crosses indicate the cluster representative, i.e., the realization closest to the centroid of each cluster.

City and consequently more scatter. This can be explained by the difference in topography. In Westport there are ridges in the north-south direction which roughly separates the onshore regions into zones that can flood independently, leading to more variation in η_{sd}^c for most realizations. However, for realizations with the most severe surface elevations, all zones are inundated, leading to smaller variation in η_{sd}^c . In Crescent City, most of the onshore region is facing the south-west direction and is rising more monotonically away from the coast, but there is a sharp gradient in the topography in the west shoreline, acting as a barrier. The variation of inundation along this barrier is significant for extreme realizations, causing more scatter in η_{sd}^c .

6 HAZARD CURVES AND MAPS

Finally we combine the techniques developed in the previous sections to produce approximate hazard curves with much less work than would be required to perform all N fine-grid simulations. We first summarize our notation and the definition of the hazard curves and these approximations. For other discussions of hazard curves, see for example (González et al., 2014), (Adams et al., 2015), and also the review paper. (Grezio et al., 2017) and associated Jupyter notebooks that illustrate these concepts interactively.

We assume that we have split all possible events into J classes indexed by $j = 1, 2, \dots, J$ (in our case $J = 4$ and the classes are for Mw 7.5, 8.0, 8.5, and 9.0). We assume each class has an associated annual probability P_j . We also make the reasonable assumption that these probabilities are sufficiently small that the probability of two events happening in a year is negligible. More specifically, we assume that the annual probability of at least one earthquake (from the classes considered) is well approximated by the sum of the P_j . In fact if the different classes represent potentially independent events, then the probability

of at least one of them occurring is given by

$$\begin{aligned}
 P_{\text{total}} &= 1 - \prod_{j=1}^J (1 - P_j) \\
 &= \sum_{j=1}^J P_j - \sum_{i \neq j} P_i P_j + \text{higher order terms} \\
 &\approx \sum_{j=1}^J P_j.
 \end{aligned} \tag{13}$$

As long as the probabilities are small, the final line of (13) is a good approximation to the true value. For the probabilities listed in Table 1, $P_{\text{total}} = 0.08527$ when calculated using the product formula in the first line of (13), and is well approximated by the slightly larger value obtained using the sum $\sum P_j = 0.08704$.

We next chose N_j realizations from each class j , for a total of $N = \sum_j N_j$ realizations. We use R_i as shorthand for “Realization i ”, for $i = 1, 2, \dots, N$ (enumerating all realizations from all of the classes). For each R_i we assign an associated *weight* w_i defined as $w_i = P_j/N_j$ if R_i is of class j .

Now consider a fixed location in the study region where we have computed h_{max} , the maximum tsunami inundation depth, for each realization. The value computed on the fine grid for realization R_i will be denoted by $h_{\text{max}}^f(R_i)$. If we perform fine-grid simulations for all realizations, then we can define the ground truth hazard curve (at this location) as follows. For any exceedance value $\hat{h} \geq 0$ we might choose, let $\{i : h_{\text{max}}^f(R_i) > \hat{h}\}$ denote the indices of the set of realizations for which h_{max} computed on the fine grid exceeds this value. Then we define

$$P[h_{\text{max}} > \hat{h}] = p^f(\hat{h}) \equiv \sum w_i \quad \text{summed over } \{i : h_{\text{max}}^f(R_i) > \hat{h}\}. \tag{14}$$

Plotting $p^f(\hat{h})$ versus \hat{h} gives the hazard curves, as shown for example in Figure 9. From the hazard curve at each point on a grid covering the study region, it is possible to extract the data needed to produce a hazard map; see Section 7.2.

Note that summing the weights w_i over all R_i for which \hat{h} is exceeded, as done in (14), is equivalent to computing $\sum_{j=1}^J (\hat{N}_j/N_j)P_j$, where \hat{N}_j is the number of realizations from Class j for which $h_{\text{max}}^f(R_i) > \hat{h}$. (Since each $w_i = P_j/N_j$ for some j and we add in one such contribution for each realization that exceeded \hat{h} .) We refer to the w_i as weights rather than probabilities because we do not mean to imply that every realization in a class has the same probability of occurring, even though they each have the same weight. Some of the realizations may be outliers that are very unlikely to occur, while most of them will come from closer to the center of the distribution. But because we assume that we sampled the distribution within each class properly, the fraction \hat{N}_j/N_j is the proper frequency to modify the probability P_j , and our choice of weights accomplishes this via the definition (14).

In Section 5 we discussed an approach to clustering the R_i into clusters indexed by $k = 1, 2, \dots, K$, for some number of clusters K that is much smaller than N . For each cluster we identified one realization \bar{R}_k from the cluster that we will call the “cluster representative”, with the hope that a single fine-grid simulation of the tsunami resulting from \bar{R}_k will give a good indication of the flooding expected for all realizations in the cluster.

One simple strategy for approximating the hazard curve is then to assign a weight \bar{w}_k to Cluster k , defined by

$$\bar{w}_k = \sum w_i \quad \text{summed over } \{i : R_i \text{ is in cluster } k\}, \quad (15)$$

and then approximate $P[h_{\max} > \hat{h}]$ by a function we will denote $p^{cf}(\hat{h})$ with the superscript standing for ‘‘cluster-fine’’:

$$p^{cf}(\hat{h}) = \sum \bar{w}_k \quad \text{summed over } \{k : h_{\max}^f(\bar{R}_k) > \hat{h}\}. \quad (16)$$

In other words, we assume that if the tsunami from the cluster representative \bar{R}_k gives an h_{\max} exceeding \hat{h} then all realizations in the cluster will, and we add in the cluster weight \bar{w}_k in this case, which is just the sum of the weights w_i for all realizations in the cluster.

Another approximation, which requires no clustering, would be to approximate $P[h_{\max} > \hat{h}]$ by

$$p^c(\hat{h}) = \sum w_i \quad \text{summed over } \{i : h_{\max}^c(R_i) > \hat{h}\} \quad (17)$$

where $h_{\max}^c(R_i)$ is the h_{\max} value obtained with a coarse grid simulation of realization R_i . This uses information from all N realizations, but is presumably not a good approximation because, by definition, the coarse grid is not sufficiently fine to resolve the study region adequately.

Much better results can be obtained by using all of the coarse-grid results after enhancing them using the techniques presented in Section 4. Using only the coarse-mod corrections to each coarse-grid result would lead to the approximation

$$p^{cm}(\hat{h}) = \sum w_i \quad \text{summed over } \{i : h_{\max}^{cm}(R_i) > \hat{h}\} \quad (18)$$

while also using the clustering to produce the pseudo-fine corrections gives

$$p^{pf}(\hat{h}) = \sum w_i \quad \text{summed over } \{i : h_{\max}^{pf}(R_i) > \hat{h}\}. \quad (19)$$

In general using this as an approximation to the true hazard curve defined by $p^f(\hat{h})$ has been found to give very good results with much less work. Only K fine-grid simulations need to be performed, but the results are based on all N realizations, with pseudo-fine approximations that are often nearly as good as the true fine-grid results when incorporated into PTHA.

7 PTHA RESULTS

We now explore the results of performing PTHA using the clustering strategies developed in Section 5, either alone or in conjunction with additional coarse-mod or pseudo-fine results as developed in Section 4.

Recall that we have sampled $N = 2000$ realizations of a CSZ event using the techniques described in Section 2, and for the purposes of this paper we assume that the hazard curves (and resulting maps) that are generated from a fine-grid tsunami simulation of each of these events is the correct reference solution, which we are trying to approximate more cheaply using the clustering and pseudo-fine grid techniques. To assess the accuracy of our approximations, we performed fine-grid simulations of all realizations in order to compute $p^f(\hat{h})$, although in practice this is what we wish to avoid.

7.1 Hazard Curves

Recall from Section 6 that a hazard curve is defined at each point in the study region where h_{\max} values have been calculated over the entire simulation. Figure 9 shows sample hazard curves at two locations

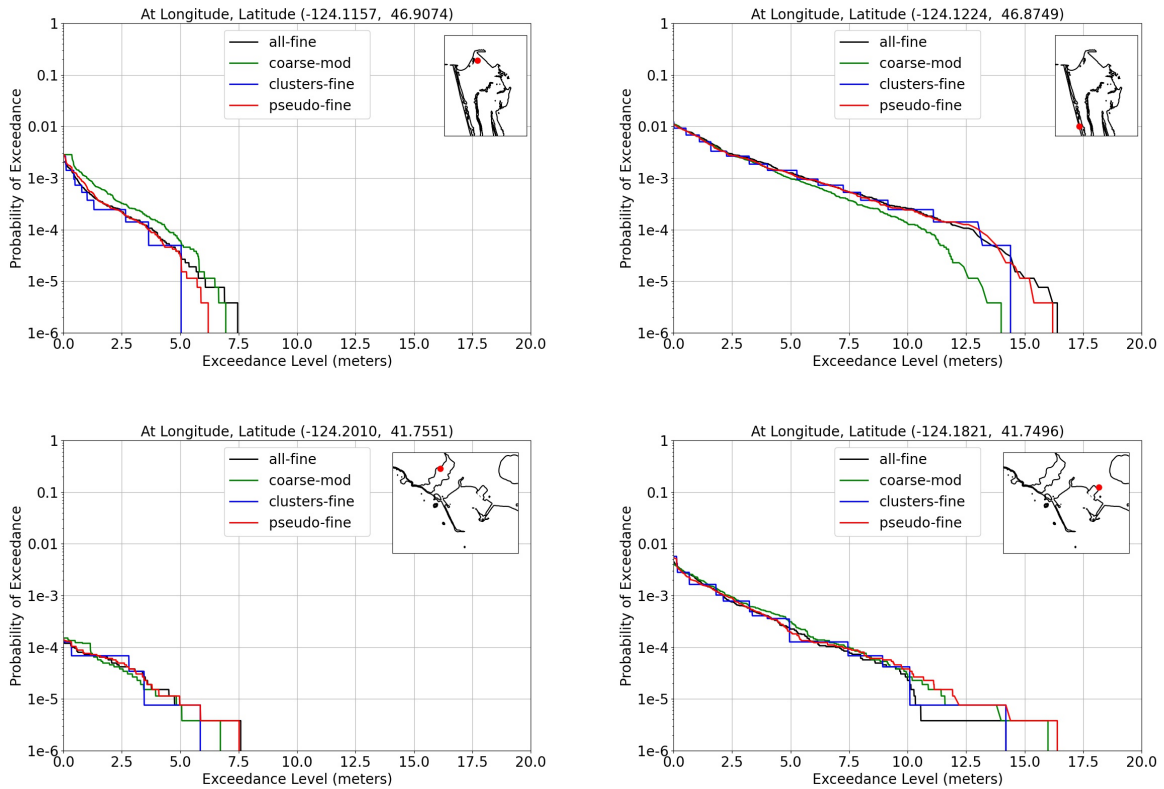


Figure 9. Sample hazard curves at two locations in Westport (top) and two locations in Crescent City (bottom), in each case indicated by the dots in the inset maps. The curves show the reference all-fine result $p^f(\hat{h})$ and three approximate hazard curves $p^{cm}(\hat{h})$, $p^{cf}(\hat{h})$, and $p^{pf}(\hat{h})$ as defined in Section 6.

in Westport, and two in Crescent City. At each location the figure shows the reference curve $p^f(\hat{h})$ and three of the approximations discussed above, the clusters-fine, coarse-mod, and pseudo-fine strategies. The particular spatial points were chosen to illustrate typical hazard curves. Additional hazard curves can be found on the website Williamson et al. (2020).

Note that the all-fine, coarse-mod, and pseudo-fine hazard curves obtained using 2000 realizations typically have 2000 jump discontinuities, one at the location of h_{\max} for each realization. The magnitude of the jump in probability at each discontinuity is equal to the weight w_i assigned to that realization. This is because this h_{\max} value contributes w_i to the estimated annual probability for any smaller exceedance value, but not for any larger exceedance value. Also note that the smallest nonzero probability that can occur on any hazard curve is the weight we assign to $M_w 9$ events, $w_4 = P_4/500 = 3.8 \times 10^{-6}$ where P_4 is from Table 1.

The hazard curve $p^{cf}(\hat{h})$ for the cluster-fine strategy is computed using only the fine-grid results $h_{\max}^f(\bar{R}_k)$ for the 18 cluster representatives. As a result, it has only 18 jump discontinuities and the jump in annual probability at each discontinuity is the cluster weight \bar{w}_k . This generally gives a reasonable approximation to the true hazard curve within the constraint of a piecewise constant function with so few jumps. It may not agree well for the most extreme events (smallest probabilities) since it assigns 0 annual probability to any exceedance value \hat{h} greater than the maximum of $h_{\max}^f(\bar{R}_k)$ over the $k = 1, 2, \dots, 18$, whereas the true hazard curve goes to 0 only above $\hat{h} = \max_i h_{\max}^f(R_i)$ maximized over all 2000 realizations. If the realization that maximizes this happens to be a cluster representative then the two hazard curves indicate

the same maximum possible flooding, but in general this will not be the case. Similarly, for other very small values of p that correspond to inundation by only a few of the 2000 realizations, the probability can not be properly represented when only using the 18 cluster representatives.

Using the coarse-mod enhancement of each coarse-grid result gives the hazard curve $p^{cm}(\hat{h})$. Even though all N simulations are now used, this correction is not sufficient to give good results in general, and this hazard curve generally deviates significantly from the correct hazard curve $p^f(\hat{h})$. However, using the pseudo-fine version of all N coarse-grid simulations gives much better results, as seen in Figure 9 and also generally seen at other locations.

In evaluating the results shown in Figure 9, it is important to remember that we cannot expect very good agreement at the smallest annual probabilities, where the results depend entirely on the most extreme tsunamis out of the 2000 selected. Of most interest in this study is the portion of the hazard curve above say $p = 10^{-4}$, which corresponds to a return time of $T = 10,000$ years. Developing an accurate hazard curve for lower probabilities would require more than 2000 realizations, even considering only the aleatoric uncertainty. Moreover, the epistemic lack of knowledge of the proper probability distribution would also be a limiting factor.

7.2 Hazard Maps and Transects

The hazard curves $p^f(\hat{h})$ at every point on the h_{\max} grid can be combined to produce hazard maps, as follows. For a fixed annual probability \hat{p} we determine from each hazard curve the corresponding value of \hat{h} such that $p^f(\hat{h}) = \hat{p}$. After doing this at every grid point in the study region we can produce plan view plots of the expected inundation depth \hat{h} for this \hat{p} . Figures 10 and Figures 11 show such plots for two different probabilities, $p = 0.002$ (return time $T = 500$ years) and $p = 0.0004$ (return time $T = 2500$ years). In each case we show the results for four of the strategies listed above. Again the all-fine strategy gives the reference result, and we compare this to the cluster-fine, coarse-mod, and pseudo-fine strategies. In general the pseudo-fine strategy gives the best approximation to the all-fine results.

Note that in these maps we show offshore points as well as onshore points, since h_{\max} for both the fine-grid and coarse-grid simulations were obtained by monitoring the maximum water depth on rectangular grids also covering some offshore points. At these points h_{\max} is always at least as great as the original pre-seismic water depth, so at these points we do not plot h_{\max} itself but rather the quantity we call zeta, defined by $\zeta_{\max} = h_{\max} + B$, where B is the pre-seismic topography at the point. More generally we define

$$\zeta_{\max} = \begin{cases} h_{\max} & \text{if } B > 0, \\ h_{\max} + B & \text{if } B \leq 0. \end{cases} \quad (20)$$

Then ζ_{\max} agrees with h_{\max} onshore, is continuous at the shoreline, and offshore it indicates the maximum tsunami elevation relative to sealevel. (We always use the pre-seismic topography in defining this, since each realization can have a different amount of uplift or subsidence.)

It is hard to quantitatively compare these hazard maps, and impossible to present results for more than one probability \hat{p} on the same map of this type. So in Figures 10 and Figures 11 we also show two selected transects in each study region (each at some fixed latitude). We then plot a cross section of the hazard map along this transect for four different values of $T = 1/p$ as indicated in the legend. As T increases the expected flood level naturally increases. In these plots we also compare two strategies, the reference all-fine result as a dashed line and the pseudo-fine strategy as the solid line. These plots clearly show that the pseudo-fine strategy does a remarkably good job of estimating the correct inundation for all four return

times, at least along these particular transects. These are fairly typical of the results seen at other locations in the study regions, and additional plots are available on the webpage [Williamson et al. \(2020\)](#).

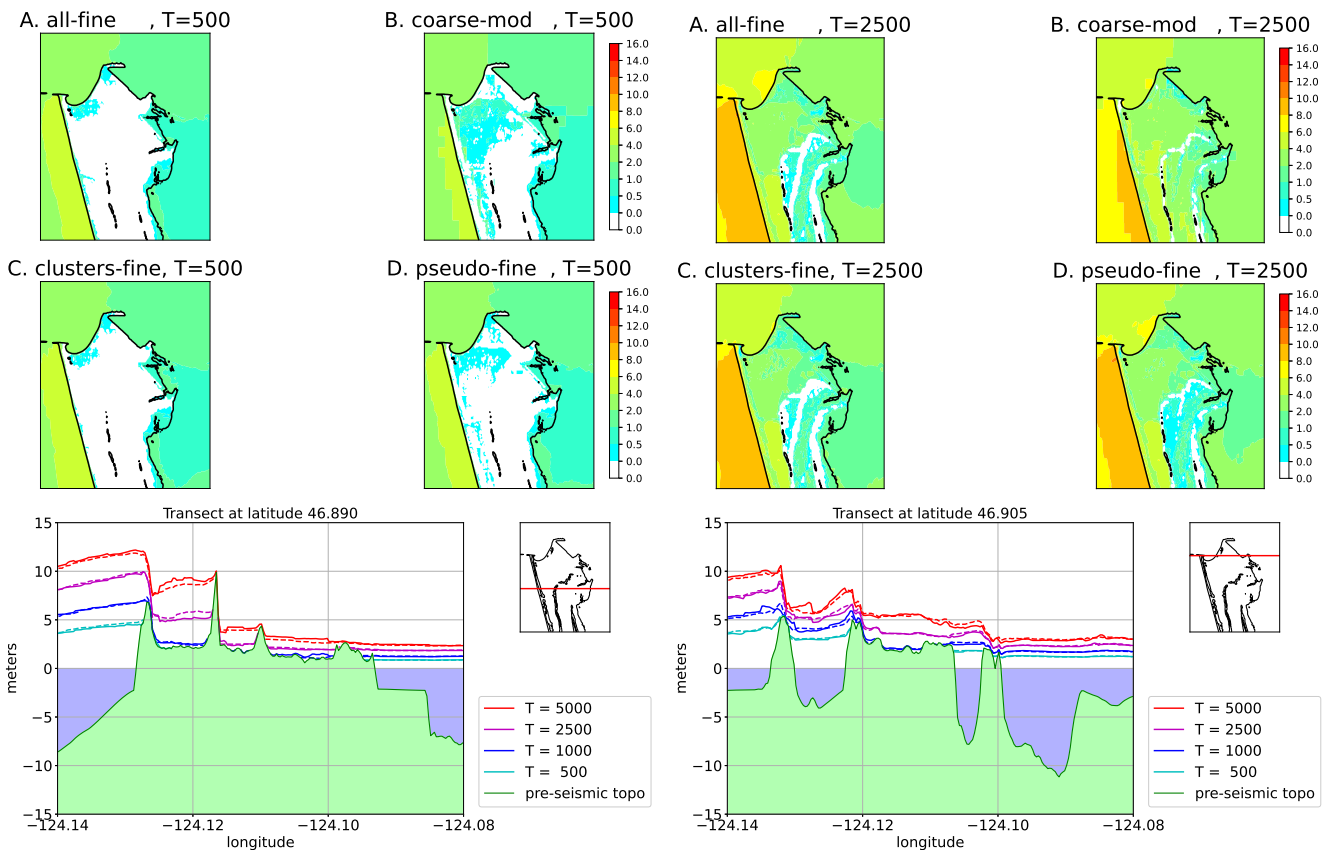


Figure 10. Sample hazard maps and transects for select locations in Westport. The top panels show plan view plots for return times $T = 500$ years (left) and 2500 years (right). In each case panel (A) shows the reference all-fine hazard maps produced with the all-fine hazard curves $p^f(\hat{h})$, (B) shows the map produced with the coarse-mod hazard curves $p^{cm}(\hat{h})$, (C) shows the map produced with the clusters-fine hazard curves $p^{cf}(\hat{h})$, (D) shows the map produced with the pseudo-fine hazard curves $p^{p^f}(\hat{h})$. The bottom figures show transects of the hazard maps for four different choices of annual probabilities p , corresponding to different return times $T = 1/p$ as indicated in the legend. In each case the dashed line is the all-fine reference solution, while the solid line is the approximation generated using 18 clusters and the pseudo-fine improvements of the other coarse-grid runs.

Finally, in Table 2, we give the maximum and mean differences between ζ_{\max} over each study region for two different return times, for each of the three strategies illustrated above. Note that in each case, the pseudo-fine strategy has mean errors less than 15 cm, even though the mean value of ζ_{\max} varied from 1.88 to 5.18 meters in the four cases shown ($T = 2500$ and 500 years, at the two study regions). This indicates that our approach is capable of giving very small relative errors, in the maximum flooding depth of the tsunami, even for the longer return time shown.

8 CONCLUSIONS

We have presented a general approach to performing PTHA when given (a) a set of classes of possible events with an annual probability or return time for each class, and (b) a probability density within each class that can be sampled to obtain a sufficiently large number of sample realizations that hazard curves

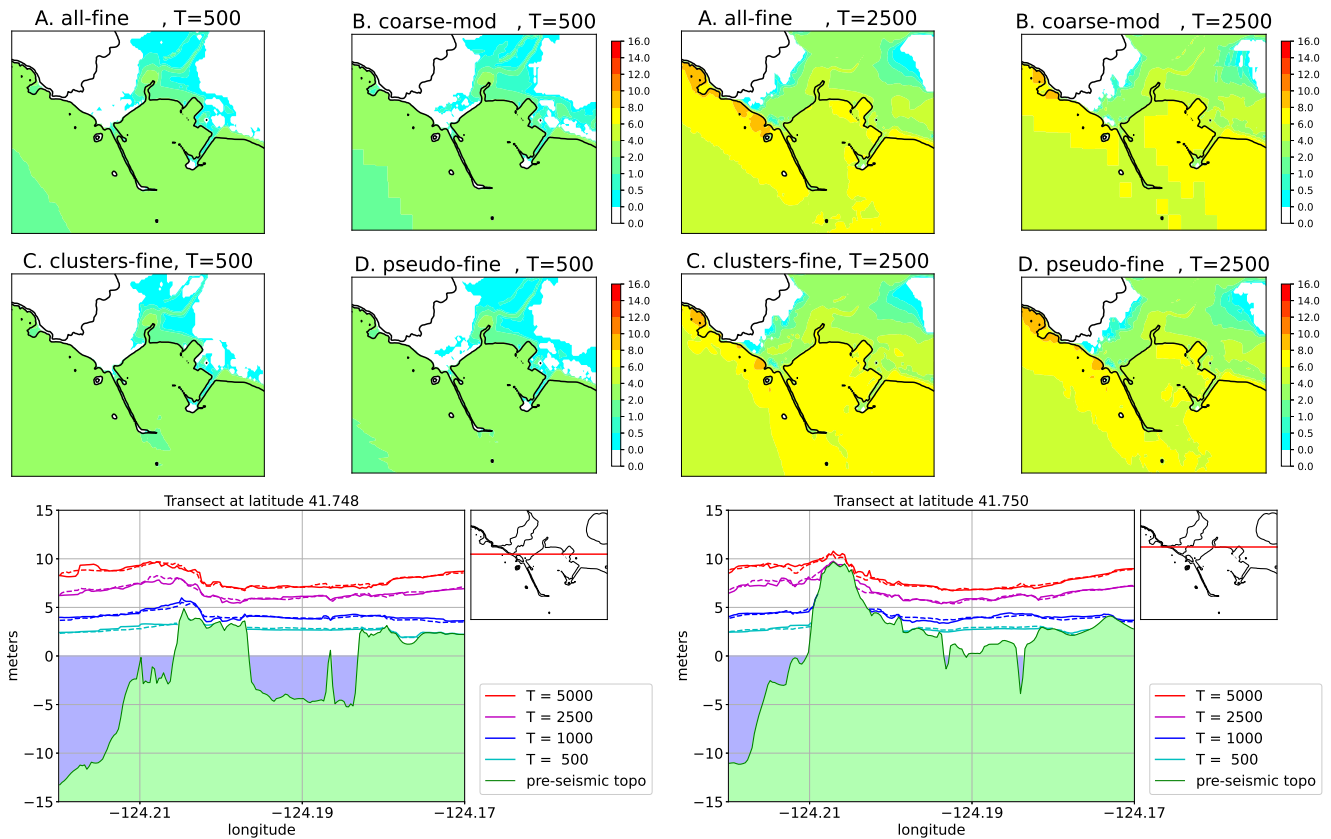


Figure 11. Sample hazard maps and transects for Crescent City. As described in the caption to Figure 10.

can be accurately approximated by performing fine grid tsunami simulations for each realization. The problem we considered is that the number of realizations needed may be too large to perform fine grid simulations of each, particularly if many coastal locations are of interest, and so the goal is to obtain good approximations to the hazard curves that would be generated by all the fine-grid simulations with much less work, employing only coarse-grid simulations, clustering, and correction procedures.

We considered a model problem where 2000 fine-grid inundation simulations were performed in order to obtain a reference hazard curve to test our methodology, which we again summarize. We first performed coarse-grid inundation simulations for each realization, using a set of four magnitude classes and a K-L expansion to define the probability density within each class for illustrative purposes. We then clustered them into only 18 clusters and performed one fine-grid simulation for a single representative from each cluster. We also used these 18 simulations, with very little additional work, to enhance the remaining coarse-grid results. The resulting 2000 pseudo-fine results were then used to produce approximate hazard curves that are much more accurate than those obtained using the cluster representatives alone.

Although we chose geophysically reasonable parameters, we do not claim that the results of our fine-grid hazard curves are correct, only that they are reasonable reference solutions against which to compare cheaper strategies. We have also ignored other magnitude events on the Cascadia Subduction Zone, other fault mechanisms such as splay faults, along with distant earthquakes and other tsunami sources, so the results in this paper should not be interpreted as providing realistic estimates of hazards in either Crescent City or Westport. Certainly any realistic PTHA meant to inform decision-making should also include sensitivity studies, particularly in light of the large epistemic uncertainty in the parameters that go into the

T	2500 years				500 years			
	max(Δ)	max(ζ)	mean(Δ)	mean(ζ)	max(Δ)	max(ζ)	mean(Δ)	mean(ζ)
Westport								
all-fine	0.00	10.1	0.00	3.75	0.00	5.06	0.00	2.01
coarse-mod	5.86	8.93	0.67	3.74	3.03	4.70	0.22	1.70
clusters-fine	1.72	9.64	0.28	3.52	0.95	4.81	0.10	2.03
all-pseudo	1.19	10.1	0.13	3.62	0.75	4.87	0.07	1.88
CC								
all-fine	0.00	9.62	0.00	5.23	0.00	3.87	0.00	2.19
coarse-mod	2.54	8.86	0.32	5.21	1.58	3.35	0.45	2.10
clusters-fine	2.78	10.5	0.47	5.34	1.10	3.68	0.22	2.30
all-pseudo	1.61	9.87	0.14	5.18	0.64	3.88	0.11	2.14

Table 2. The magnitude of ζ_{\max} and differences between ζ_{\max} as computed using the coarse-mod, clusters-fine, and all-pseudo strategies, compared to the reference all-fine strategy at both Westport and Crescent City (CC), for return times $T = 2500$ and 500 years ($p = 0.0004$ and 0.002). All values are in meters. The columns labelled $\max(\Delta)$ and $\text{mean}(\Delta)$ are the maximum and mean of the difference over all grid points (i, j) , e.g. for the coarse-mod row, $\Delta = |\zeta_{\max}^{cm} - \zeta_{\max}^f|$. For comparison, the corresponding maximum and mean values values of ζ_{\max} across the community are also listed for each strategy, in the columns labelled $\max(\zeta)$ and $\text{mean}(\zeta)$. Recall that ζ represents the maximum flooding depth on land, or the flooding depth added to the pre-seismic bathymetry for points offshore.

probability densities (whether generated by K-L expansion or by any other technique). The techniques of this paper can be useful for such sensitivity studies since they can accelerate the PTHA for any set of realizations and thus allow testing more sets of realizations in order to investigate the resulting variation in hazard curves. Sets of realizations might be generated with different density parameters, or be of different sizes. Different random sets of realizations of the same size can also help to better explore the aleatoric uncertainty.

More work is needed to better understand and optimize the clustering method used in this paper. In particular there is a need to better quantify the number of clusters that should be used, and the best set of proxy variables to use in performing the clustering.

However, we note that we were generally able to achieve nearly the same hazard curves with our coarse-grid results as when using all the fine-grid results, down to an annual probability of $p = 10^{-4}$ or less for each of our test communities. We therefore believe these techniques can be a useful component in a full probabilistic study of these sites, and many others. We also note that they can be applied to any other choice of classes and probability densities, and adapted to work with any tsunami modeling software.

CONFLICT OF INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

All authors contributed to the initial development and project goals of this study. A.W. and D.M. developed the earthquake realizations and performed the tsunami simulations. F.G., R.L., and D.R. developed the clustering methodology, and L.A. developed the coarse-mod and pseudo-fine techniques. All authors contributed to writing the text and producing the figures presented.

FUNDING

This research was supported in part by FEMA Region IX, Grant EMW-2014-CA-00218, and by NASA DISASTERS grant 80NSSC19K1104, AFOSR CoE FA9550-17-1-0195, AFOSR MURI FA9550-15-1-0038, and NSF Hazard SEES grant EAR-1331412.

DATA AVAILABILITY STATEMENT

The datasets generated for this study can be found on the website ([Williamson et al., 2020](#)), along with many additional figures.

REFERENCES

- Adams, L., LeVeque, R., and González, F. (2015). The pattern-method for incorporating tidal uncertainty into probabilistic tsunami hazard assessment (PTHA). *Natural Hazards* 76, 19–39
- [Dataset] Adams, L. M., LeVeque, R. J., Rim, D., and Gonzalez, F. I. (2017). Probabilistic Source Selection for the Cascadia Subduction Zone. Results from a study supported by FEMA Region IX. Webpage: <http://depts.washington.edu/ptha/FEMA>
- [Dataset] Amante, C. and Eakins, B. W. (2009). ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis. NOAA Technical Memorandum NESDIS NGDC-24, <http://www.ngdc.noaa.gov/mgg/global/global.html>
- Blaser, L., Krüger, F., Ohrnberger, M., and Scherbaum, F. (2010). Scaling relations of earthquake source parameter estimates with special focus on subduction environment. *Bulletin of the Seismological Society of America* 100, 2914–2926
- [Dataset] Clawpack Development Team (2020). Clawpack software. Version 5.7.0
- Comninou, M. and Dundurs, J. (1975). The angular dislocation in a half space. *Journal of Elasticity* 5, 203–216
- Crempien, J. G., Urrutia, A., Benavente, R., and Cienfuegos, R. (2020). Effects of earthquake spatial slip correlation on variability of tsunami potential energy and intensities. *Scientific reports* 10, 1–10
- Frankel, A., Chen, R., Petersen, M., Moschetti, M., and Sherrod, B. (2015). 2014 update of the Pacific Northwest portion of the US National Seismic Hazard Maps. *Earthquake Spectra* 31, S131–S148
- Geist, E. and Lynett, P. (2014). Source processes for the probabilistic assessment of tsunami hazards. *Oceanography* 27, 86–93
- Geist, E. L. and Parsons, T. (2006). Probabilistic analysis of tsunami hazards. *Natural Hazards* 37, 277–314
- Geuzaine, C. and Remacle, J.-F. (2009). Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *International journal for numerical methods in engineering* 79, 1309–1331
- Goda, K., Yasuda, T., Mori, N., and Maruyama, T. (2016). New scaling relationships of earthquake source parameters for stochastic tsunami simulation. *Coastal Engineering Journal* 58, 1650010–1
- Goldberg, D. E. and Melgar, D. (2020). Generation and validation of broadband synthetic P waves in semistochastic models of large earthquakes. *Bulletin of the Seismological Society of America*
- Goldfinger, C., Nelson, C. H., Morey, A. E., Johnson, J. E., Patton, J. R., Karabanov, E. B., et al. (2012). *Turbidite event history—Methods and implications for Holocene paleoseismicity of the Cascadia subduction zone*. Tech. rep., US Geological Survey
- [Dataset] González, F., LeVeque, R., Adams, L., Goldfinger, C., Priest, G., and Wang, K. (2014). Probabilistic Tsunami Hazard Assessment (PTHA) for Crescent City, CA Final Report September 11, 2014. Univ. Washington ResearchWorks Archive: <http://hdl.handle.net/1773/25916>

- [Dataset] González, F. I., LeVeque, R. J., and Adams, L. M. (2013). Tsunami Hazard Assessment of the Ocosta School Site in Westport, WA. Univ. Washington ResearchWorks Archive: <http://hdl.handle.net/1773/24054>
- Gonzalez, J., Gonzalez, G., Aranguiz, R., Melgar, D., Zamora, N., Shrivastava, M. N., et al. (2020). A hybrid deterministic and stochastic approach for tsunami hazard assessment in Iquique, Chile. *Natural Hazards* 100, 231–254
- Grezio, A., Babeyko, A., et al. (2017). Probabilistic tsunami hazard analysis: Multiple sources and global applications. *Reviews of Geophysics* 55, 1158–1198. Supplementary material: <https://doi.org/10.5281/zenodo.816290>
- Horrillo, J., Knight, W., and Kowalik, Z. (2008). Kuril Islands tsunami of November 2006: 2. Impact at Crescent City by local enhancement. *Journal of Geophysical Research: Oceans* 113
- LeVeque, R. J., Waagan, K., Gonzalez, F. I., Rim, D., and Lin, G. (2016). Generating random earthquake events for probabilistic tsunami hazard assessment. *Pure Appl. Geophys.*
- Lloyd, S. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28, 129–137
- Mai, P. M. and Beroza, G. C. (2000). Source scaling properties from finite-fault-rupture models. *Bulletin of the Seismological Society of America* 90, 604–615
- Mai, P. M. and Beroza, G. C. (2002). A spatial random field model to characterize complexity in earthquake slip. *Journal of Geophysical Research: Solid Earth* 107, 10–1
- McCaffrey, R., Qamar, A. I., King, R. W., Wells, R., Khazaradze, G., Williams, C. A., et al. (2007). Fault locking, block rotation and crustal deformation in the Pacific Northwest. *Geophysical Journal International* 169, 1315–1340
- McCrary, P. A., Blair, J. L., Waldhauser, F., and Oppenheimer, D. H. (2012). Juan de Fuca slab geometry and its relation to Wadati-Benioff zone seismicity. *J. Geophys. Res.* 117, B09306
- Melgar, D. and Hayes, G. P. (2019). The correlation lengths and hypocentral positions of great earthquakes. *Bulletin of the Seismological Society of America* 109, 2582–2593
- Melgar, D., LeVeque, R. J., Dreger, D. S., and Allen, R. M. (2016). Kinematic rupture scenarios and synthetic displacement data: An example application to the Cascadia Subduction Zone. *J. Geophys. Res. – Solid Earth* 121, 6658–6674
- Melgar, D., Williamson, A. L., and Salazar-monroy, E. F. (2019). Differences between heterogeneous and homogeneous slip in regional tsunami hazards modelling. *Geophysical Journal International*, 553–562
- [Dataset] NOAA NCEI (Accessed 2012). Crescent City, California 1/3 Arc-second MHW Coastal Digital Elevation Model. <https://www.ngdc.noaa.gov/metadata/page?xml=NOAA/NESDIS/NGDC/MGG/DEM/iso/xml/693.xml&view=getDataView&header=none>
- [Dataset] NOAA NCEI (Accessed 2017). Astoria, Oregon 1/3 Arc-second MHW Coastal Digital Elevation Model. <https://www.ngdc.noaa.gov/metadata/page?xml=NOAA/NESDIS/NGDC/MGG/DEM/iso/xml/5490.xml&view=getDataView&header=none>
- Okada, Y. (1985). Surface deformation due to shear and tensile faults in a half-space. *Bulletin of the Seismological Society of America* 75, 1135–1154
- Parsons, T., Console, R., Falcone, G., Murru, M., and Yamashina, K. (2012). Comparison of characteristic and Gutenberg–Richter models for time-dependent $m \geq 7.9$ earthquake probability in the Nankai-Tokai subduction zone, Japan. *Geophysical Journal International* 190, 1673–1688
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830

- Satake, K. (1987). Inversion of tsunami waveforms for the estimation of a fault heterogeneity: Method and numerical experiments. *Journal of Physics of the Earth* 35, 241–254
- Sepulveda, I., Liu, P. L. F., Grigoriu, M., and Pritchard, M. (2017). Tsunami hazard assessments with consideration of uncertain earthquake slip distribution and location. *Journal of Geophysical Research-Solid Earth* 122, 7252–7271
- Volpe, M., Lorito, S., Selva, J., Tonini, R., Romano, F., and Brizuela, B. (2019). From regional to local SPTHA: efficient computation of probabilistic tsunami inundation maps addressing near-field sources. *Natural Hazards and Earth System Sciences* 19, 455–469
- Williamson, A., Melgar, D., and Rim, D. (2019). The effect of earthquake kinematics on tsunami propagation. *Journal of Geophysical Research: Solid Earth* 124, 11639–11650
- [Dataset] Williamson, A. L., Rim, D., LeVeque, L. M. A. R. J., Melgar, D., and González, F. I. (2020). Website for this paper with additional results. <http://depts.washington.edu/ptha/frontiers2020a/>