Probabilistic Source Selection for the Cascadia Subduction Zone

Final Report

March 19, 2017

Loyce M. Adams, Randall J. LeVeque, Donsub Rim, and Frank I. González

University of Washington

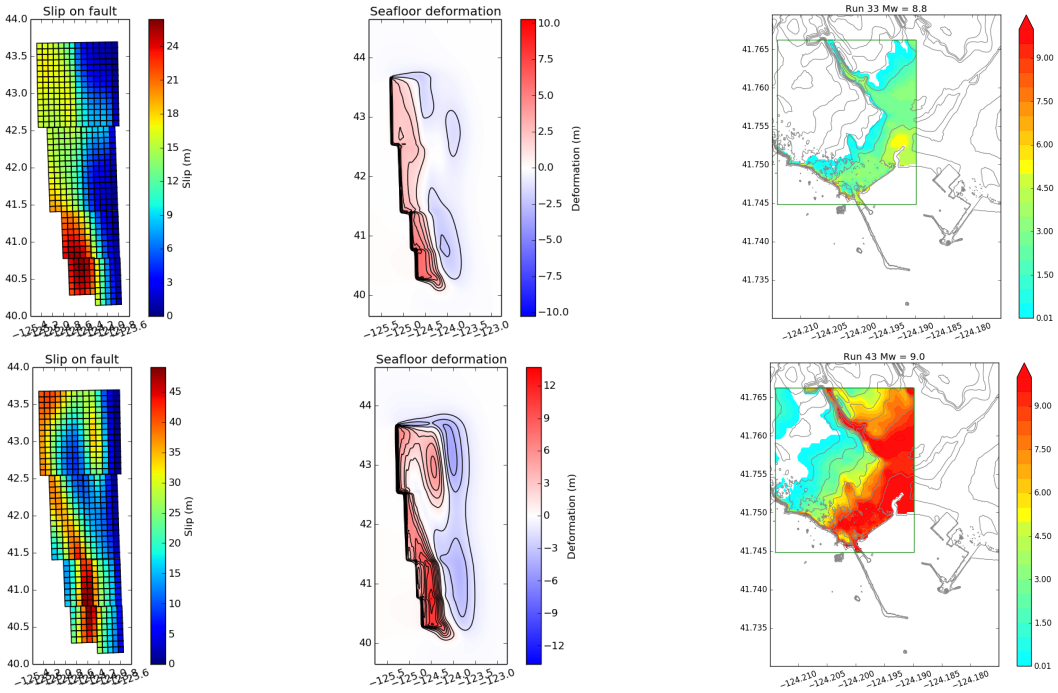

Study funded by FEMA Region IX

Contents

1 Introduction

This report has been submitted to FEMA Region IX as a final project report for a project on developing new methodologies for Probabilistic Tsunami Hazard Assessment (PTHA). We propose a methodology for taking a large number of realizations of potential future earthquakes (with associated probabilities) and producing good approximations to the resulting hazard curves and maps without doing a computationally-expensive fine-grid tsunami simulation for each realization. The specific goals and an outline of the strategy are summarized in Section 2.

We believe that this project illustrates the potential of the methodology we propose, but there is still much work to be done to optimize it and to apply it to specific hazard assessment projects. In particular, the input for this methodology is a large number of realizations/probabilities that are assumed to accurately reflect the true probability distribution of future events. For the work in this report we generated a set of 400 realizations as test data but we know they are not a good reflection of reality, as discussed further in Section 3. Generating a more realistic set was not the goal of this project.

On the methodology side there is also much that can be done to improve on the methods presented below. We have learned a great deal from the work presented here and in retrospect would do many things differently. We have listed several observations and ideas for improving the methodology in this report, and will be using it as the basis for future work on this topic.

For the tsunami simulations presented in this report we have used the GeoClaw software [3, 18], from Version 5.3.1 of the open source software Clawpack [6]. The techniques developed here are independent of what software is used for the tsunami simulation, however.

This report contains a number of figures, but many more can be viewed online from
http://depts.washington.edu/ptha/FEMA/
All of the computer code used in this work has also been archived there, including the code to generate the earthquake realizations, to run the tsunami simulations, to perform the filtering and clustering, and to produce the hazard curves and maps, and to perform comparisons of different approaches.

1.1 Background on PTHA

In 2006, Geist and Parsons [11] first adapted the 1968 Cornell Probabilistic Seismic Hazard Analysis (PSHA) methodology [7] to Probabilistic Tsunami Hazard Assessment (PTHA), primarily by substituting numerical model computations of tsunami wave height for empirical seismic attenuation relationships to determine ground motions. Since then, there has been a steady growth in research to improve the PTHA methodology and to apply it to hazard assessments on the scale of individual communities (González, et.al. [13] in 2009; Thio, et al. [28] in 2012; González et.al. [14] in 2014) and larger geographical regions (Annaka et al. [2] in 2007; Burbidge, et al. [5] in 2008; Heidarzadeh and Kijko [16] in 2011; Thio et al. [28] in 2012).

Seismic source specification is a fundamental input for both PSHA, to generate seismic waves that induce ground motion, and PTHA, to generate tsunami waves that result in coastal inundation. Similarly, specification of the wave propagation media is required for each methodology i.e., the distribution of the earth's density and elasticity for seismic waves, and the distribution of topography (water depth and land height) for tsunamis. We here assume topography is adequately known, in the form of a digital elevation model (DEM), and focus on the source specification problem in PTHA.

Most generally, for seismic tsunami generation, the temporospatial evolution of vertical and horizontal crustal deformation are required as the initial and forcing conditions for tsunami modeling, at the appropriate time and length scales. Okada [23] provides frequently-used closed expressions for seismic deformation due to a point source, given the depth and the dip, strike and tensile components of the slip vector, $\mathbf{s}$, and, in the case of a finite rectangular fault plane, the additional parameters of length, L, and width, W. Then, for extended finite sources such as those that might occur on the Cascadia Subduction Zone (CSZ) many such point or relatively small fault plane sources must be distributed on the surface of a finite fault model, such as the slab model in Blair, et.al, [4], or, more recently, the model of [9]. At this point, then, the seismic source

specification problem becomes that of assigning the temporospatial distribution of slip vectors, $\mathbf{s}$, over the model fault plane surface.

The temporal dependence of $\mathbf{s}$ does have an effect on the generation and evolution of a tsunami, but this effect is typically small compared to spatial variations in $\mathbf{s}$ because the seismic rupture speed is substantially larger than the tsunami propagation speed; Geist [10] provides a review of theoretical, numerical and experimental investigations into this issue. Also see, in particular, the analysis of Keller [17] and the theoretical and experimental work of Hammack [15]. For simplicity, therefore, we here ignore any time-dependency of $\mathbf{s}$ and focus only on its spatial variations; i.e., we make the frequently adopted assumption of instantaneous crustal deformation that produces an initial tsunami waveform that is an exact copy of the ocean bottom displacement.

Suites of possible earthquake scenarios for a given region can be developed through expert geological and seismological knowledge that takes account of field measurements that include geodetic data and paleoseismic evidence, and geophysical constraints such as fault geometry and slip budgets; a logic tree is frequently employed to organize and categorize earthquakes by recurrence interval, fault mechanism, magnitude, length, width, and other seismic parameters. As an example, the National Seismic Hazard Mapping program of the U.S. Geological Survey periodically updates and publishes PSHA products for the nation (Petersen, et al. [25]; Petersen, et al. [26]). Similarly, in 2013 Witter, et al. [29] developed suites of earthquake scenarios that provide the spatial resolution for the slip distributions required to conduct PTHA studies; these scenarios were subsequently used to assess the impact of partial CSZ ruptures on the Washington and northern Oregon coasts in 2014 by Priest et al. [27] and in the 2014 PTHA study of Crescent City by González et.al. [14]. For the work presented here we generated random earthquake realizations using an approach described briefly in Section 3 and in more detail in Appendix G.

1.2 Hazard curves and hazard maps

We assume that the general PTHA methodology of [14] will be used to construct hazard maps for a region of interest by first construcing hazard curves at each point on a grid covering this spatial domain. The hazard curve at each point shows a quantity of interest (such as the maximum depth of flooding) on the horizontal axis and the probability of exceeding this value on the vertical axis (e.g. see Figure 10). In [14] these were generated by running a fine-grid tsunami simulation for each possible realization and combining these using their associated probabilities.

In that work we also incorporated tidal uncertainty by running simulations at different possible tide stages and incorporating information about the tidal patterns at Crescent City, using methods we also published in [1]. Here we ignore tidal uncertainty, although the methods we develop could easily be extended to that case at the cost of additional simulations at different tide stages (which would make the efficiency of the filtering techniques developed here even more important).

Once a hazard curve is available at each point on a spatial grid, it is then possible to produce a hazard map showing the depth of flooding expected at each point for some particular fixed probability. (Or showing the probabilities associated with a fixed flooding depth, an alternative product we advocate in [14] may be useful.)

The goal of the current project is to efficiently compute good approximations to the hazard curves at each spatial point.

1.3 The study region in Crescent City, CA

We chose a rectangular region of Crescent City, CA as our study region, the region enclosed by the blue rectangle in Figure 1. We wanted a relatively small region containing an interesting variety of topography. We specified an 88×78 grid of points at 2/3 arc-second resolution where the hazard curves will be computed.

Figure 1 also shows two transects (in yellow) and below we show some of the results along these transects. The western transect was chosen to go through downtown Crescent City while the eastern transect goes up

a narrow valley at the northen end. This valley is particularly interesting since the water depth may reach high values in this valley when simulations are peformed on a fine grid, while the resolution of the coarser grids used in this study do not resolve the valley at all.

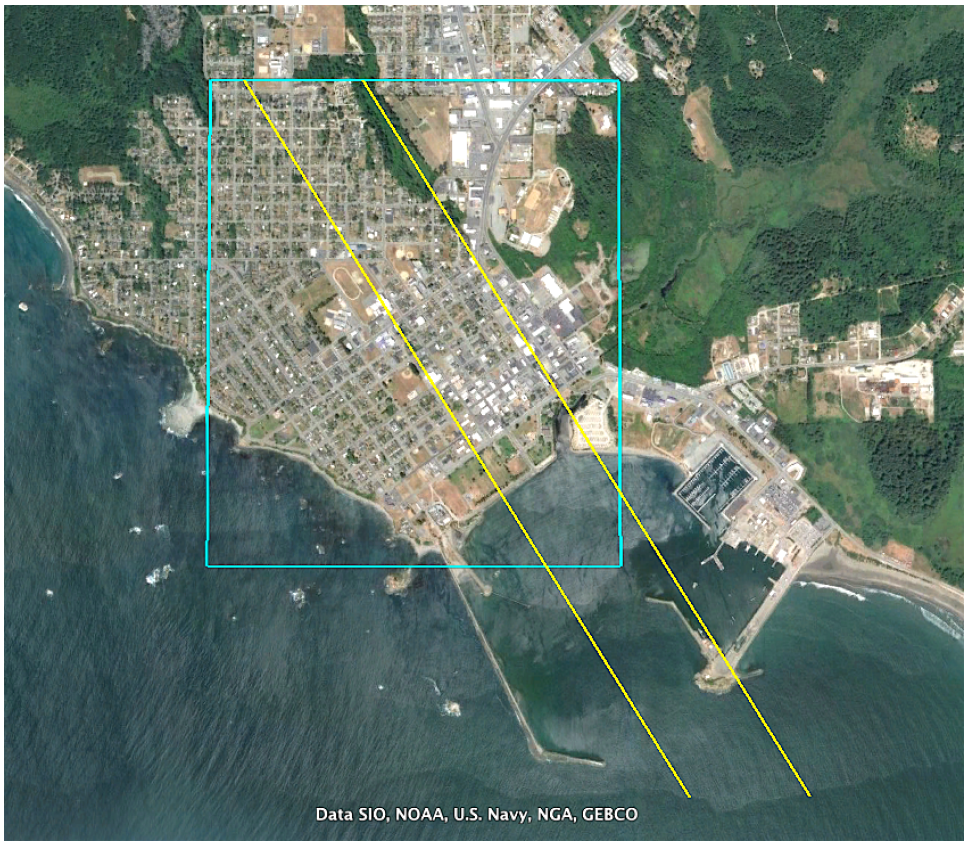

Figure 1: The rectangular region of Crescent City, CA that was used for this study (blue box) and two transects (yellow lines).

2 Overview

Here we summarize the problem that was the primary focus of this study and the approaches we have taken to solving this problem and to assessing and comparing different methods.

1. **Problem:**
   (a) **Given:**
      i. Some large number N of hypothetical earthquake events with associated probabilities (e.g $N = 400$ in this study). We assume these probabilities sum to 1 (i.e. we consider conditional probabilities on possible realizations of an event with unspecified return time).
      ii. A specific region where hazard maps are needed. Here we use a portion of Crescent City, CA as described above.
   (b) **Produce:**
      i. Hazard curves at each point on a grid of points covering the region of interest. A hazard curve shows the probability of exceedance as a function of flooding depth, and can be constructed from computational results showing the maximum depth of flooding at this spatial point for each of the hypothetical events. (Note that these would be "conditional hazard" curves based on assuming the events considered have probabilities summing to 1. These would have to be weighted by the probability of the overall event, based on its return time, and combined with similar curves obtained by considering other possible tsunami sources in order to obtain the full hazard curves.)
      ii. Hazard maps (e.g. the 500-year flood map, which shows the depth of flooding associated with any fixed probability, or maps that show the probability of flooding for a fixed depth). These can be produced from the hazard curves at each point using methods described elsewhere, e.g. [14].
2. **Brute force approach:** Simulate all N events on a sufficiently fine computational grid (using e.g. GeoClaw) in order to capture the maximum flooding depth at each point on the regional grid. We assume this gives the "exact" hazard curve at each point but is too expensive.
3. **Goal of this study:** Develop a methodology to produce good approximations to the hazard curves in a way that requires far fewer GeoClaw simulations (e.g. $M = 20$ rather than $N = 400$).
4. **Strategy:**
   (a) Perform cheap computations for each event to obtain one or more *proxy* values that are correlated with the severity/pattern of flooding. In particular we have explored:
      i. Performing GeoClaw simulations for all N events but on a much coarser grid (e.g. 10 arc-second rather than 2/3 arc-second) since coarse grid simulations can be run much more quickly.
      ii. Calculating proxies based only on the seafloor deformation of each earthquake event. This is typically the way the events are described and so this is very cheap, but gives proxies that are not as well correlated with the quantities of interest.
   (b) Cluster the events into $M \ll N$ clusters of "similar" events based on the proxy values.
   (c) Run a single fine-grid simulation for one representative event from each cluster.
   (d) Produce hazard curves based on these M fine grid simulations by either:
      i. Assigning a probability to each of the M representative events by summing the probabilities of all events in its cluster, or
      ii. Using the fine grid solutions from the representative event to improve the coarse grid solutions for other events in the cluster (giving *pseudo-fine grid* results) and then using all N events to construct the hazard curve (based on the pseudo-fine grid results for events that do not have fine-grid results).

5. **Testing and comparing strategies:**

   (a) We generated $N = 400$ "random" earthquake events on the southern portion of the Cascadia Subduction Zone (CSZ) and assigned a probability to each, as described in Section 3.

   (b) We used the brute force approach described above to compute the "exact" hazard curve at each point on this grid, based on GeoClaw simulations of each event with full 2/3 arc-second resolution (about 20 m) in the region around Crescent City (adaptive mesh refinement (AMR) is used, so the grid is coarser elsewhere). This can then be used to assess the accuracy of various approximation approaches.

   (c) We also ran a coarse-grid GeoClaw simulation with a resolution of 10 arc-seconds (about 300 m) around Crescent City for each of the $N = 400$ events.

   (d) Although the coarse-grid results are lacking many details of the flow, we found that several proxy quantities determined from these results are highly correlated with similar values from the fine-grid runs. This is explored in Section 4.

   (e) We developed and tested several strategies for clustering and for using the resulting small set of fine-grid runs (and perhaps pseudo-fine runs) to construct hazard curves and maps. The "exact" results were used to assess and compare the strategies.

3 Earthquake realizations used in this study

The original terms of this contract specified that several hundred earthquake events would be provided to us as a test suite, but those never materialized. Instead we have developed our own approach to generating random earthquakes that has recently been published [19, 21]. We use a Karhunen-Loève (K-L) expansion to generate a random slip pattern on a specified fault geometry, sampled from random fields with a specified covariance structure. Using this K-L Expansion, we developed a methodology that provides a catalogue of earthquake scenarios of desired magnitude along with the corresponding annual probabilities. For the purposes of this study, we generated 400 realizations using the approach described briefly here and in more detail in Appendix G.

Note: We do not claim that the 400 realizations we used in this study are sampled from the correct probability distribution of possible future events, in fact we know they are not. The problem of tuning the fault geometry and covariance functions used in the K-L expansion to produce a realistic probability distribution is the subject of ongoing work. For the purposes of this study, however, it does not matter whether the realizations or probabilities are realistic since we are studying a methodology for taking a set of N samples with given probabilities and producing good approximations to the resulting hazard curves with minimal work. The "exact" solution for the set of realizations used is known (as computed via N fine grid simulations as described above), and we believe that our findings based on the set of realizations we used should carry over to other sets of realizations.

With this in mind, we chose the realizations to give a broad range of impacts at Crescent City. We generated slip on a southern portion of the Cascadia Subduction Zone (shown in Figure 50) in a way that tended to give highly concentrated patches of slip, and we varied the magnitude from M_w 8.6 to 9.2. The largest events are clearly non-physical, with concentrated patches of slip exceeding 180 m in some cases (roughly 3 times more than observed in the extreme Tohoku event of 2011). However, since this slip is quite concentrated the resulting tsunami is often no worse or in some cases much smaller than what is computed with uniform slip at the same magnitude (but sometimes much worse). We consider the great variation in computed results to be a good test of our methodology and expect that our methods would work better if we assumed a more realistic variation in slip patterns.

We created 100 slip patterns using the approach outlined in Appendix G. Each of these was then scaled to 4 different magnitudes, M_w 8.6, 8.8, 9.0, and 9.2, using the fact that multiplying the slip on each subfault (and hence the seafloor deformation) by approximately a factor of 2 increases the moment magnitude by 0.2. In this manner we generated a total of $N = 400$ realizations, 100 realizations at each magnitude M_w 8.6, 8.8, 9.0, and 9.2. We could have generated an independent set of 100 realizations at each magnitude rather than scaling up a smaller set, but since past work on PTHA (e.g. [14]) has used the "Bandon sources", which consist of only 3 slip patterns scaled up to 5 magnitudes each, we decided to follow a similar approach with our much larger set of realizations.

Each of the 100 slip patterns had a probability between 0 and 1 assigned to it, with these probabilities summing to 1. The method used to assign probabilities is also described in Appendix G. To assign probabilities to the full set of 400 realizations, we weighted the probability of each slip pattern by weights of 0.3, 0.3, 0.3, and 0.1 for each of the 4 realizations that it generated (of magnitude 8.6, 8.8, 9.0, and 9.2, respectively). The smaller weight for the largest class is loosely based on the relatively small weight assigned to the XXL events in the collection of Bandon sources. Again we stress that these probabilities need not be realistic in order to test different methodologies in the context of this study.

Some of the approaches we investigate work best if only a few of the realizations account for a large fraction of the probability distribution (as an extreme example, if 20 realizations have probability summing to 1 and the other 380 have probability 0, then clearly one only needs to do 20 fine grid simulations to compute the exact hazard curves). So as another test of methods, we have also considered the same set of realizations but with the probability set to 0.01 for each of the 100 original slip patterns. This uniform distribution of probabilities gives a more difficult problem if one wants to estimate hazard curves accurately from only $M = 20$ fine grid simulations.

Figure 2 shows a few of the 400 realizations used in this study. A range of realizations has been chosen to illustrate the variety. The full set of 400 can be viewed at http://depts.washington.edu/ptha/FEMA/.

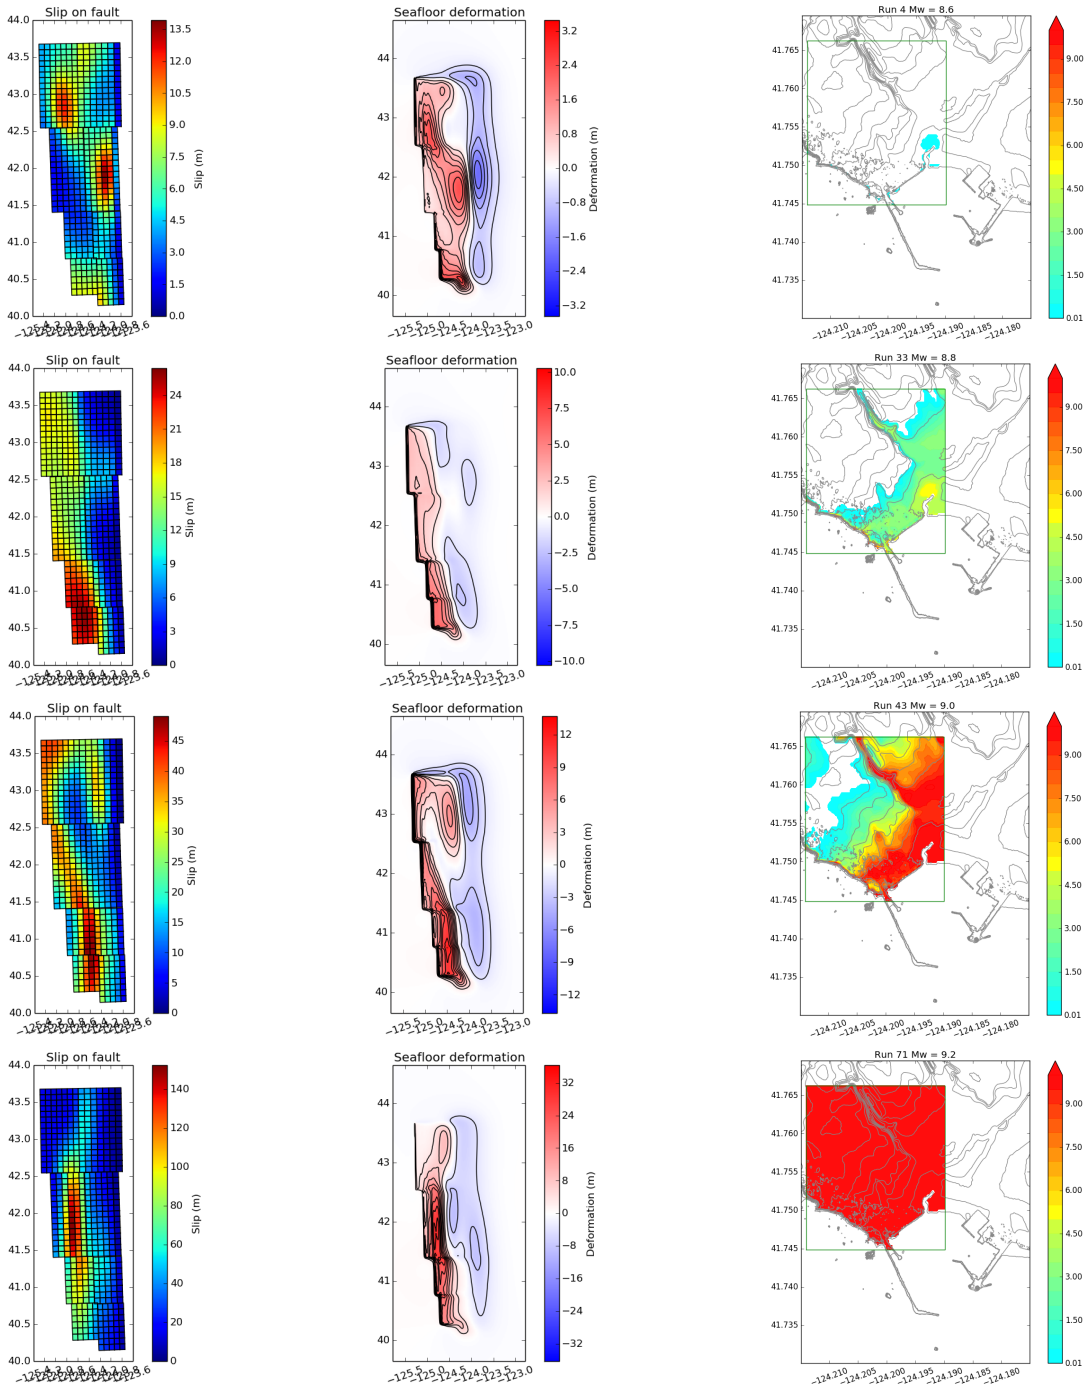

Figure 2: Four sample realizations. The left column shows the slip distribution on the CSZ fault geometry, the middle column shows the resulting sea floor deformation, and the right column is the maximum flooding depth in Crescent City.

4 Comparison of computed results for the 400 realizations

In this section we show some of the computed results from both fine-grid simulations and the coarse-grid (cheaper) simulations that were used in this study in order to further illustrate the range of inundation patterns that were observed, and the degree to which the coarse-grid simulations give a reasonable approximation.

4.1 Pseudo-color plots over the region

The bottom row in Figure 3 shows a comparison of the maximum inundation depth, as computed on the fine grid and the coarse grid, for one sample realization (run 5, magnitude 8.6). The top row in the figure gives the absolute value of the inundation ($h_{\max}$) differences and the absolute value of the eta (η) differences. This shows that the coarse grid alone is inadequate to represent the inundation, and illustrates the degree of coarsening (a factor of 15 in each coordinate direction). Nonetheless we will see that information from the coarse grid runs can be used to good advantage, after they have been cleaned as described in Section 7.3 and Appendix A.

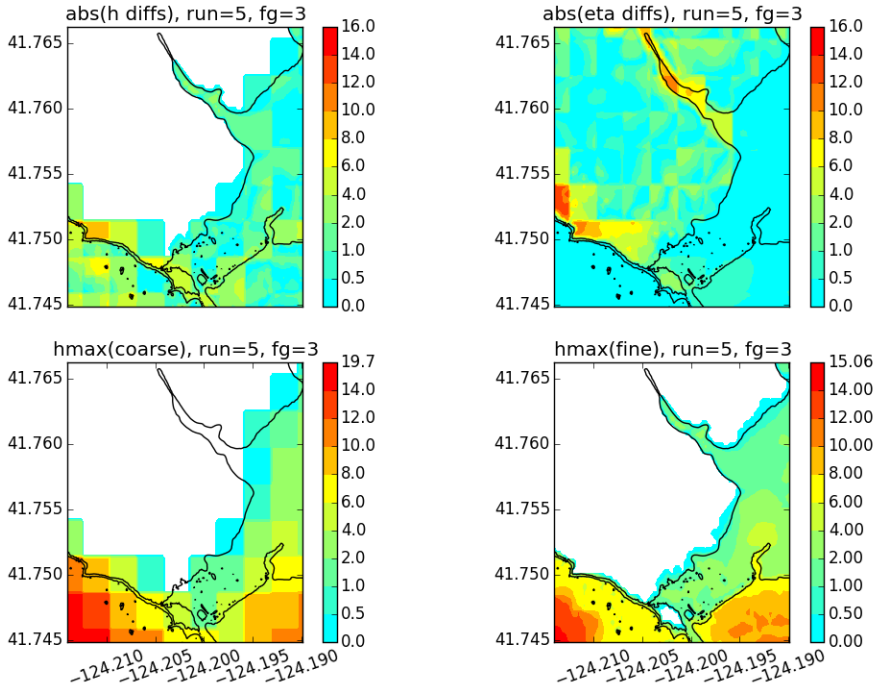

Figure 3: Mw 8.6, run 5: The bottom row shows the coarse and the fine data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the coarse and fine data.

4.2 Transect plots

The pcolor plots shown in Figure 3 are difficult to compare across many realizations and also do not show the elevation of the flood above sea level directly, since the color is based on flooding depth.

It is also useful to plot the maximum elevation of flooding along a transect. The plots in Figure 4 show these plots for 50 of the events from each magnitude class (the other 50 are similar but the plots become

unreadable with more lines).

What is actually plotted is the value of η at each point along the transect, where $\eta = h + B$ with B the topography and h the maximum flow depth recorded over the entire simulation.

Note that the earthquake generally causes uplift or subsidence at Crescent City, and these plots show η relative to the initial topography.

Note the following from the transect plots:

- For each event the value of η varies along the transect but is much more nearly constant than the flow depth h is. A long-wave tsunami moving in slowly tends to flood the entire flooded region to roughly uniform elevation, and the depth at each point then depends on the topography. If $\eta(x, y) = \eta_0$ everywhere for some event, then the depth would be simply $h(x, y) = \eta_0 - B(x, y)$. This is an important observation discussed further below.

- These plots show how coarse the coarse grid is relative to the fine grid (300 m vs. 20 m resolution). The maximum flood depth is stored on a fine grid along the transect in either case, but this fine grid cuts through coarse grid cells in the coarse grid calculation and so a step function is observed for the topography on the coarse grid. (The length of each step is not 300 m because the transect cuts through different lengths of each computational cell since it is at an angle to the rectangular computational grid.)

- In spite of the coarseness of the coarse grid, there is a high degree of correlation between the results on the two grids. This is explored more below and exploited in our methodology.

- Transect 2 shows much more discrepancy between fine and coarse results than is seen on Transect 1, particularly at the north end. This transect was chosen to run up a narrow valley (see Figure 1) that is well resolved on the fine grid but completely unresolved on the coarse grid.

- There is a great deal of variation within each magnitude class and if they were all plotted together on a single plot they would be seen to overlap substantially. There is even a Mw 9.2 realization that exhibits less flooding than most of the Mw 8.6 events (this particular realization had slip concentrated in such a way that there was extreme uplift at Crescent City and much of the surface displacement occurred onshore rather than generating a tsunami). This indicates that simply running a few realizations from each magnitude class might not be the best approach to clustering.

From the fine grid runs:

From the coarse grid runs:

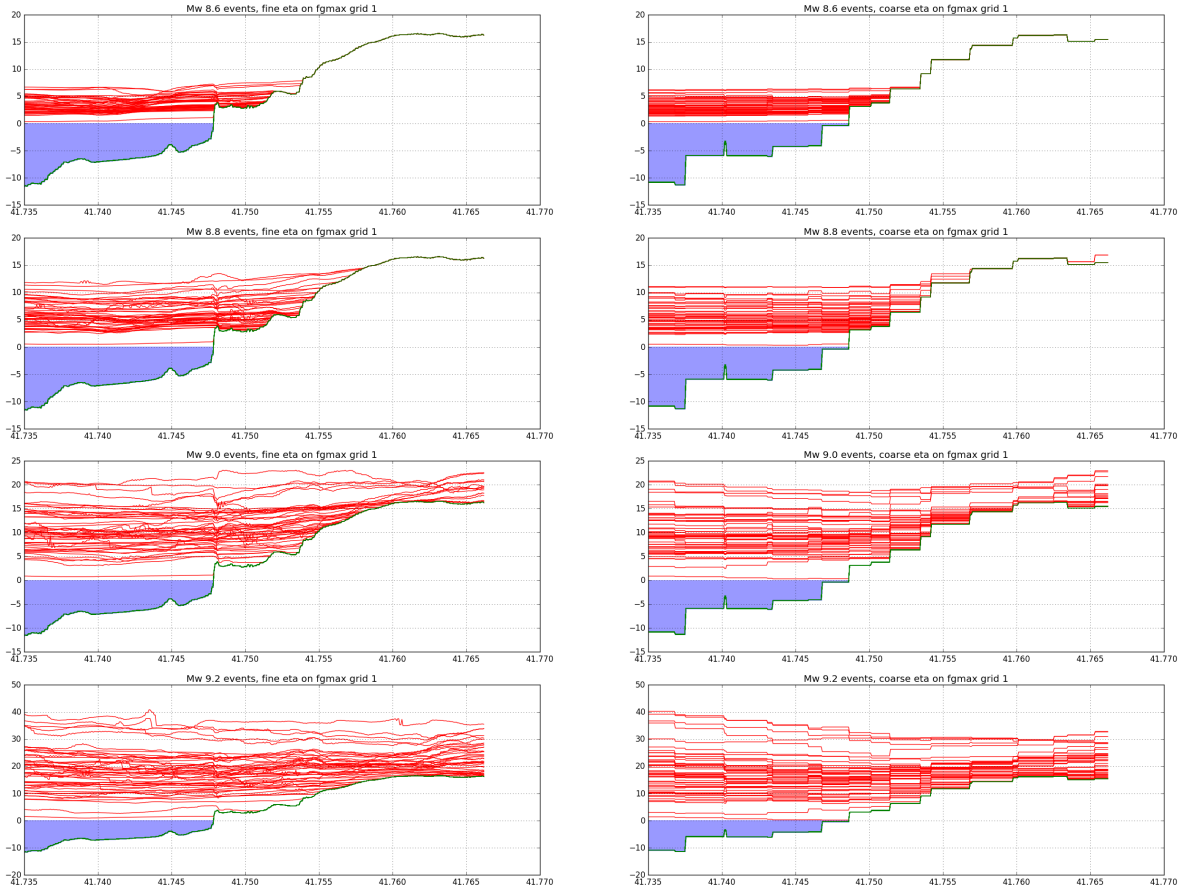

Figure 4: Maximum η for 50 realizations from each magnitude class along Transect 1. The left column shows results computed on the fine 2/3 arc-second grid, the right column shows results from the same 50 realizations as computed on the coarse 10 arc-second grid. In each case the rows correspond to Mw 8.6, 8.8, 9.0, and 9.2.

From the fine grid runs:

From the coarse grid runs:

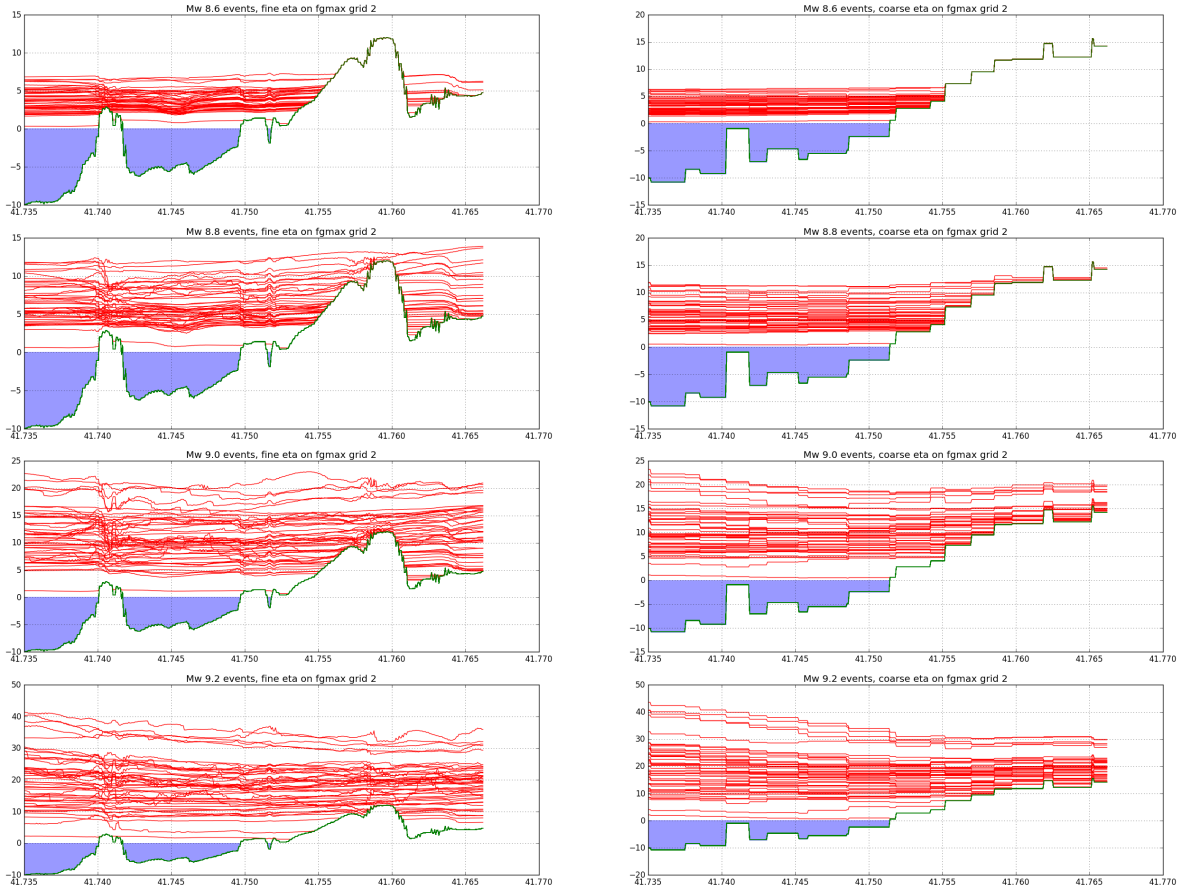

Figure 5: Maximum η for 50 realizations from each magnitude class along Transect 2. The left column shows results computed on the fine 2/3 arc-second grid, the right column shows results from the same 50 realizations as computed on the coarse 10 arc-second grid. In each case the rows correspond to Mw 8.6, 8.8, 9.0, and 9.2.

4.3 Quantities of Interest

Ultimately we are interested in building hazard curves that show the exceedance probability vs. depth at each point on the grid covering the region of interest, and for this we need to take into account the maximum flooding depth at each point from each event. The behavior of this maximum depth varies dramatically not only between different events but also between different points. To get an overall comparison of events over the full region of interest, it is useful to define some quantities associated with each event. These are used for comparing fine and coarse grid simulations, for investigating correlations between different quantities, and for performing clustering of realizations as needed for step 4(b) of the strategy outlined in Section 2. This clustering is discussed in detail in Section 5 below.

The following quantities are based on GeoClaw simulations:

- EtaMax: η_{max} for each realization is the maximum value of η over all on-shore locations where any flooding occurs. We also use η^f_{max} and η^c_{max} to denote the values computed from the fine grid run and the coarse grid run, respectively.
- EtaMean: η_{mean} for each realization is the mean value of η over all on-shore locations where any flooding occurs. We also use η^f_{mean} and η^c_{mean} to denote the values computed from the fine grid run and the coarse grid run, respectively.
- Onshore volume at Crescent City: V_{cc} for each realization is obtained by summing the maximum flow depth at each point on the grid and multiplying by the area represented by each point. Since the points are on a grid with 2/3 arc-second resolution, each point represents an area of approximately $20 \times 14 = 280m^2$ at the latitude of Crescent City. This volume represents the total volume of water that would cover the region of the grid shown in Figure 1 if all of the maximal flow depths occurred at the same instant in time. In practice they do not, so the total mass of water varies with time and is never this large. But this volume gives another good proxy for the severity of an event since it takes into account both the area covered by the flood and the maximum depth observed at each point.

The following quantities can be determined very cheaply for each realization without any GeoClaw simulation:

- Uplift/Subsidence at Crescent City: ΔB_{cc} is the vertical displacement at the location of Crescent City due to the sea floor deformation of a particular realization.
- EtaMax for initial sea surface: η^{sea}_{max} is the maximum uplift of the sea floor (and hence of the sea surface at the initial time in a GeoClaw simulation, where it is assumed that the water column is instantaneously lifted as the sea floor moves).
- Potential Energy: E_{pot} is the potential energy of the initial displacement of the sea surface.
- Displaced Volume: V_{disp} is the total volume of water displaced by the initial displacement of the sea surface.

4.4 Scatter plots of fine vs. coarse results

In the plots shown in Figure 6, each dot represents a single realization and plots the value of some quantity of interest from the fine grid run vs. the same quantity as computed on the coarse grid run (for the same realization, i.e. using the same `dtopo` file in GeoClaw, which specifies the sea floor deformation). The points are colored by the magnitude of the event.

Plots on the left show the full range of values observed, while plots on the right are zoomed in near the origin to better show the relationship for the smaller magnitude events. The first row shows the quantity η_{mean}, the second row η_{max}, and the third row shows V_{cc}.

Note the following:

- In general larger values are observed in results from the fine grid run than from the coarse grid simulation of the same realization (i.e. the points tend to lie above the 1:1 line). This is not surprising since the fine grid topography better captures extremes.
- The quantity η_{mean} is better correlated between fine and coarse grid runs than is η_{max}. This is also not surprising since on the fine grid the maximum water depth often occurs in the narrow valley, for example, which is not well resolved on the coarse grid. The mean value of η is less sensitive to these points.
- The generally good agreement between fine and coarse grid values of η_{mean} shows that the coarse grid simulation, although missing most of the details of the flow, is still capturing the general behavior of the tsunami well. This gives hope that we can use the cheaper coarse grid simulations results for clustering events as discussed in Section 5.

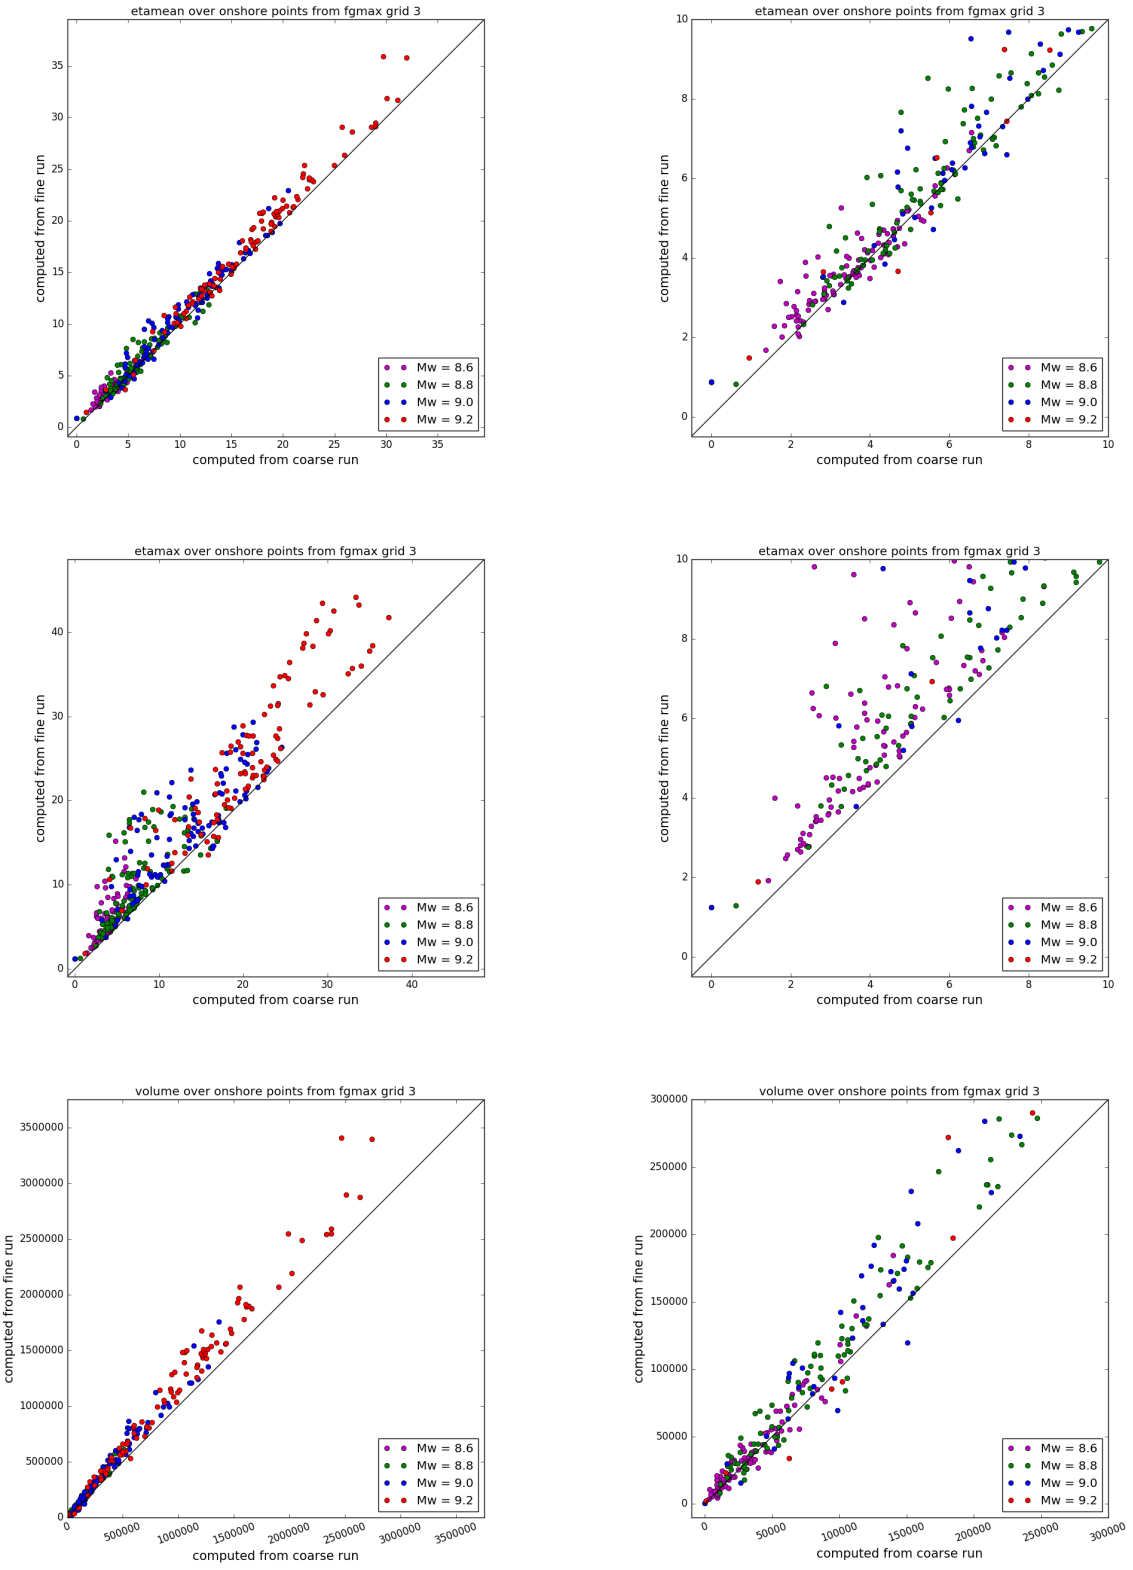

Figure 6: Scatter plots of values obtained from the coarse grid run for each realization, vs. the value of the same quantity obtained from the fine grid run for the same realization. The left column shows all 400 realizations, the right plot is zoomed near the origin to better show the smaller events. The top row shows scatter plots for η_{mean}, the middle row for η_{max}, and the bottom row for the displaced volume V_{cc}.

5 Clustering techniques

In order to reduce the number of fine grid runs, we look for ways of clustering the realizations into some small number of clusters, say $M = 20$ clusters, with each of the $N = 400$ realizations falling in one cluster. The hope is that doing a fine grid run for one realization in each cluster will give results that are representative of all realizations in that cluster.

Clustering by magnitude M_w alone is not sufficient since the 400 realizations are sampled from only four magnitudes. Moreover it is clearly seen from the plots above that there is great variation within each magnitude and a great deal of overlap of results from different magnitudes.

We consider two techniques for clustering, although there are many other possibilities and we consider this a ripe area for future research. The first is based on the use of information obtained from coarse grid runs of all N realizations, assuming the computing budget allows this. The scatter plots of Section 4.4 indicate that even a grid that is 15 times coarser in each direction can provide a good indication of the severity of the event, as judged by the proxy values considered there. This is described in Section 5.1.

The second approach is based only on proxy values that can be computed directly from the sea floor deformation of each realization and so does not require any GeoClaw runs for the clustering, only the M fine grid runs, one for each cluster. This is described in Section 5.2.

5.1 Eta mean / Eta max clustering

In this section we assume we have done all $N = 400$ coarse grid runs and can use proxy values obtained from these runs to do the clustering. One simple approach would be to cluster based on η_{mean} alone, e.g. order the events based on their η_{mean} value and then split them into clusters. However, we attempt to do a better job of clustering by using information also about η_{max} for each realization.

Figure 7 shows scatter plots of η_{mean} vs. η_{max} for each realization. Unlike the scatter plots shown above, in this case we are comparing these quantities for each realization with the same resolution grid (not comparing fine results vs. coarse results). The plot on the left shows η_{mean} vs. η_{max} computed on the fine grid for each realization, while the plot on the right shows η_{mean} vs. η_{max} computed on the coarse grid for each realization. In each case it must be that $\eta_{max} \geq \eta_{mean}$ and so all points lie on or above the 1:1 line in these plots. The further above the line, the more discrepancy there is between the maximum η observed over the entire grid and the mean value.

As expected, there is more discrepancy for larger tsunamis than for small ones, and also more discrepancy seen on the fine grid simulations (that better represent the detailed topography) than on the coarse grids. We also observe, however, that the coarse grid is reasonably effective at displaying the range of possibilities. Finally, we observe that for fixed value of η_{mean} there are often many events with very similar η_{mean} but quite different values of η_{max}. We posit that these correspond to tsunamis that behave differently and perhaps flood different parts of the region even though they give the same value of η_{mean}.

We cluster based on the location of the realizations in this $\eta_{mean} - \eta_{max}$ plane. Figure 8 shows an example where we have specified 20 clusters on the top, or 40 clusters on the bottom. Now the colors indicate the realizations that are grouped together into a single cluster. The + symbol in each cluster shows the centroid and the circle that is larger than the other circles in each cluster shows the realization that has been chosen to be representative of the cluster. (Note that we do not always choose the realization that is closest to the centroid, instead we choose the one that is closest to a weighted average of the realizations, weighted by their probabilities.)

The clustering is done using the standard procedure of k-means clustering, grouping the realizations into the specified number of clusters so that each sum of the squares of the distances from all points to the center of the clusters they belong to is minimized. The clustering is performed using the `KMeans` function from the Python module `sklearn.cluster`[1].

[1] `http://scikit-learn.org/stable/modules/generated/sklearn.cluster.KMeans.html`

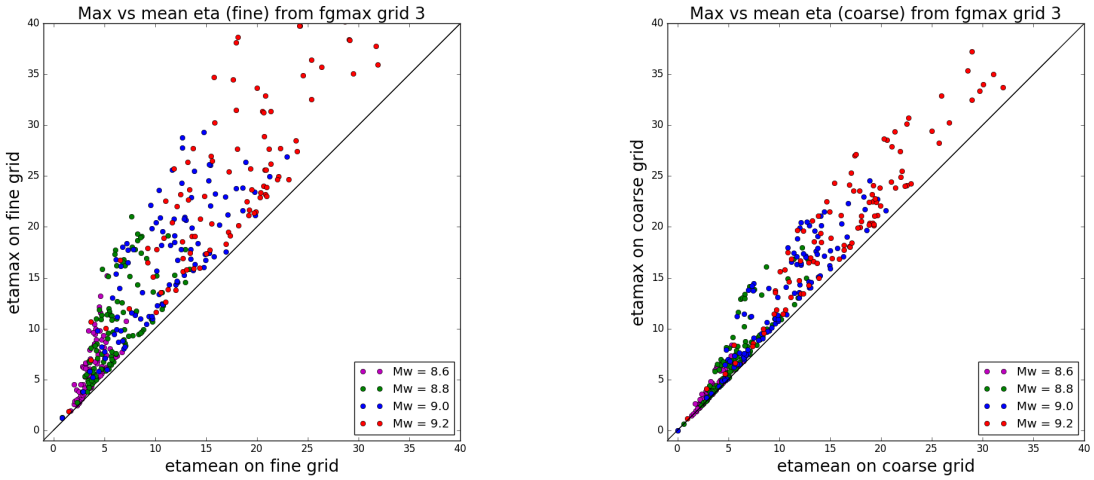

Figure 7: Scatter plots of η_{mean} on the horizontal axis vs. η_{max} on the vertical. Each dot corresponds to one of the 400 realizations. The plot on the left shows values obtained from the fine grid runs, while the plot on the right shows results for the coarse grid runs.

In Appendix H.1 we show the maximum flooding for at most 5 realizations in each cluster, which gives some indication that the clustering is working well by clustering together realizations that give similar levels of flooding. All 400 realizations arranged by clusters can be viewed at http://depts.washington.edu/ptha/FEMA/.

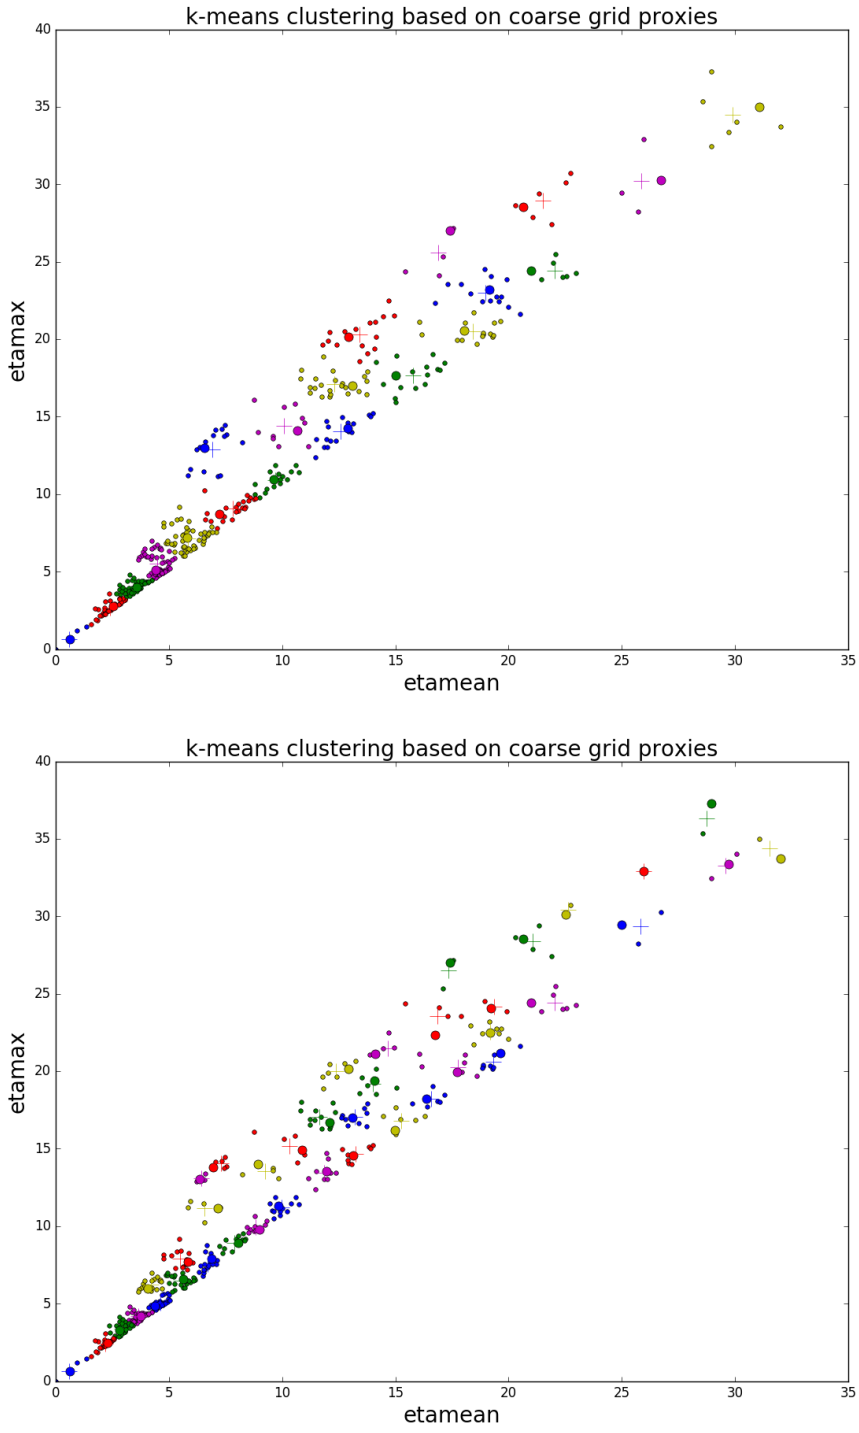

Figure 8: Clustering of realizations based on their location in the $\eta_{mean} - \eta_{max}$ plane, as determined from coarse grid simulations. In the top figure 20 clusters are indicated by different colors, in the bottom figure there are 40 clusters. The larger dot in each cluster marks the realization for which a fine-grid simulation is performed. The + marks the centroid of the cluster.

5.2 Dtopo clustering

In the previous section we presented an approach to clustering based on proxy values computed from the coarse grid simulations. We now consider the possibility of clustering using proxy values that are computed directly from the sea floor deformations of the realizations. This requires no tsunami simulations and is very cheap to perform. This approach would extend more easily to a situation where the number of realizations N is much larger than 400. We call this "dtopo clustering" since in GeoClaw the files specifying the sea floor deformation (the changes in topography) are called dtopo files.

In Section 4.3 we presented several possible proxy quantities. Here we consider just one possible clustering strategy based on the quantities ΔB_{cc} (uplift or subsidence at Crescent City) and η^{sea}_{max}, the maximum initial surface elevation offshore. We experiemented with others as well, but this case worked reasonably well and is sufficient to illustrate the potential of using cheap proxies.

Figure 9 shows 20 and 40 clusters obtained by clustering based on these quantities. Since clustering is performed based on Euclidean distance in the plane, we scaled η^{sea}_{max} by 3 so that the variations were comparable in the two coordinate directions.

Note that with the fault geometry we assumed, many of the realizations give uplift ($\Delta B_{cc} > 0$) and quite extreme values (± 15 m) are observed due to the unphysically extreme M_w 9.2 realizations included in our test set.

In Appendix H.2 we show the maximum flooding for at most 5 realizations in each cluster, which shows that this clustering strategy does not work as well as the previous more expensive strategy of Section 5.1. All 400 realizations arranged by clusters can be viewed at http://depts.washington.edu/ptha/FEMA/.

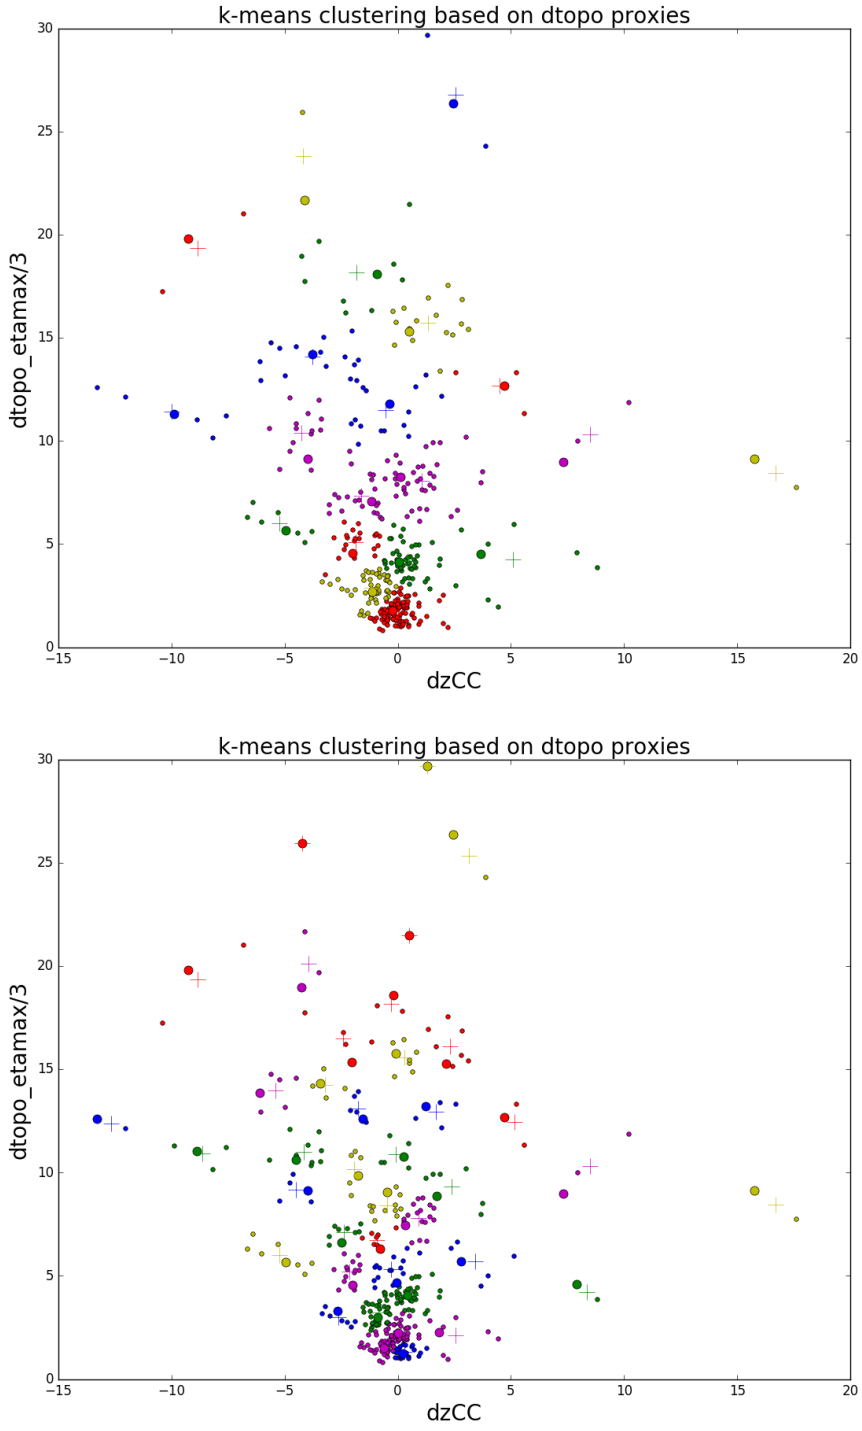

Figure 9: Clustering of realizations based on their location in the $\Delta B_{cc} - \eta^{sea}_{max}$ plane. In the top figure 20 clusters are indicated by different colors, in the bottom figure there are 40 clusters. The larger dot in each cluster marks the realization for which a fine-grid simulation is performed. The + marks the centroid of the cluster.

5.3 Comments and possible improvements

The two clustering approaches we have presented above were chosen to give simple illustrations of how this general methodology might be applied, either based on results from coarse grid simulations of every realization or using only information that is available before doing any simulations.

We show in the next few sections that even these simple clustering strategies can be very effective when used to compute approximate hazard curves and maps using relatively few fine grid simulations, only one per cluster.

But there are many improvements that could be made in this methodology and in the future we intend to explore other approaches to clustering. In particular, working with the data and results presented here has led to the following observations and ideas for future work.

1. With the wide range of severity of events used in this study, there is much more scatter in all of the proxy values at the high end than at the low end. As a result, the majority of the clusters are devoted to the extreme events and each of these clusters has few members (in some cases only 1 or two realizations when 40 clusters were specified). In contrast the clusters at the lower end of severity were very large, with dozens of realizations. So the variation that exists within the less severe events is not as well sampled. In practice one might want to impose a maximum cluster size or use some other mechanism to optimally spread out the given budget of clusters (determined by the budget of fine-grid simulations possible).

2. We have illustrated clustering using only two proxy values to facilitate showing plots of these clusters in the plane. One could cluster in a higher-dimensional space, e.g using η_{mean}, η_{max}, V_{cc} in the case of clustering based on coarse-grid runs, or using all 4 proxy values listed in Section 4.3 for dtopo. We have not yet experimented with how best to do this, in particular each proxy would have to be scaled appropriately for clustering based on Euclidean distance to make sense.

3. The clustering approach is a special case of the classification problem that has been extensively studied in Statistics and more recently under the guise of machine learning. We have chosen a simple approach to explore the potential of this methodology, but many more sophisticated methods have been developed that should be explored in the future.

4. Our strategy requires choosing one realization from each cluster to use for the fine grid simulation. In this study we chose the realization that is closest to the weighted mean of all realizations in the cluster, weighted by the probability assigned to each realization. The rationale was that if there are only two events in a cluster and one is much more likely than the other, then we should perform the fine grid simulation on the more likely one (so that we are using the correct inundation for the one that most affects the hazard curve). But for clusters that contain more than two events, the one closest to the weighted centroid may still have very small probability compared to others in the cluster. In retrospect, it may have been better to choose the event from each cluster that has the highest probability rather than using centroid information at all (except perhaps if all events have equal or very similar probabilities).

6 Probability products for the fgmax grids

The study area in Crescent City is shown in Figure 1. We collect GeoClaw data on two fgmax grids called Transect 1 and Transect 2 and on one fgmax rectangular grid. Transect 1 is the leftmost yellow line in Figure 1 and has 439 grid points. Transect 2 is the rightmost yellow line and also has 439 grid points. The third fgmax grid is the blue rectangle and it has 88x78 grid points. All grid points are in longitude and latitude.

This study assumes a Cascadia earthquake occurs due to only one of the 400 sources, and the conditional probability of each of the sources is given. These 400 conditional probabilities sum to 1. We consider two different ways of assigning these probabilities. First, the generation of the 400 sources from the KL-expansion gives a natural assignment of the probability of occurrence for each of the sources as described in Section 3. Second, for comparison, we consider a uniform distribution within each Mw category (still using the probabilities 0.3, 0.3, 0.3, and 0.1 for Mw 8.6, 8.8, 9.0, and 9.2 realizations, respectively.) Hence each of the 400 sources will have conditional probability 0.003 or 0.001 of occurrence when a Cascadia earthquake occurs.

These 400 sources were used for 400 fine grid and 400 coarse grid GeoClaw runs to propagate the tsunami into Crescent City where the maximum value of water depth, $h_{\max}$, seen during the tsunami was recorded on a 2/3 arc sec fgmax grid. This grid covered the blue rectangle shown in Figure 1. We also collected $h_{\max}$ at locations lying on two different transects shown as the yellow lines in Figure 1. The fine grid runs actually calculated $h_{\max}$ at this 2/3 arc sec resolution in Crescent City. The coarse grid runs calculated on a much coarser grid (15 fine cells per coarse cell for the transects and 225 fine cells per coarse cell for the blue rectangular region), but we recorded the results at the same 2/3 arc sec resolution using piecewise constant interpolation with all points in the same coarse cell receiving the same $h_{\max}$ value.

From the $h_{\max}$ values and the conditional probabilities, various products can be calculated, including hazard curves at individual locations, probability contours for exceeding fixed inundation levels, and contours of the maximum inundation for a given probability. For the two transects in Figure 1, the contours reduce to plots of probability for exceeding fixed inundation levels or plots of the maximum inundation for a given probability versus latitude along the transect.

In particular, we find probability contours of exceeding 0.0 to 12.0 meters of water and ζ contours for probability 0.1 to 0.9 in increments of 0.1, and for 0.95. We also consider another product called the *weighted difference* in the ζ contours by weighting across all the considered probabilities. This is done by multiplying the ζ difference for a given p by p, summing across all the p values and dividing by their sum (.1+.2 + ... + .9 + .95 = 5.4) to get a weighted difference. Of course, other weightings could be considered, depending on which values of p are of interest.

Examples of some of these products will be given in Section 7 when we describe our approximation strategies and in Sections 8, 9, and 10 when we compare them. In particular, these sections use one hazard curve at a fixed location, the probability contour for ζ exceeding 2 meters, the ζ contour for $p = 0.1$ and the weighted difference ζ contour as described above.

7 Approximation strategies

Below we outline strategies for approximating the results of the 400 fine grid runs. In Sections 7.1 - 7.5 we assume we have a budget of 20 fine grid runs that we are allowed to do and that we can afford to do all 400 GeoClaw runs on a coarse grid. The remaining 380 fine grid runs were done just for the purpose of evaluating our strategies – their results were not used in the development of the strategies. The SVD strategy described in Section 7.6 uses a budget of 32 fine grid runs and a modification of the remaining 368 coarse grid runs, and we compare it to strategies that use a budget of 40 fine grid runs.

7.1 Highest 20 probability strategy

Since the conditional probabilities of all 400 sources are known a priori, we could choose the 20 that have the highest probability of occurrence. This seems an obvious choice, and indeed it works quite well when the distribution of the conditional probabilities is such that the highest 20 represent a large fraction of the probability. For the KL probabilities the highest 20 represented 42.8% of all the probability. However, for the uniform probabilities, most of the sources have the same probability, and depending on the sort routine used to return the highest 20, the probability space will not be represented well. Also, all 9.2 Mw sources would be left out of the top 20. The left plot in Figure 10 shows that the hazard curves of all 400 fine and the highest 20 strategy for the location indicated by the red dot compares fairly well when the KL probabilities are used, but the right plot shows it does not compare well when the uniform probabilities are used.

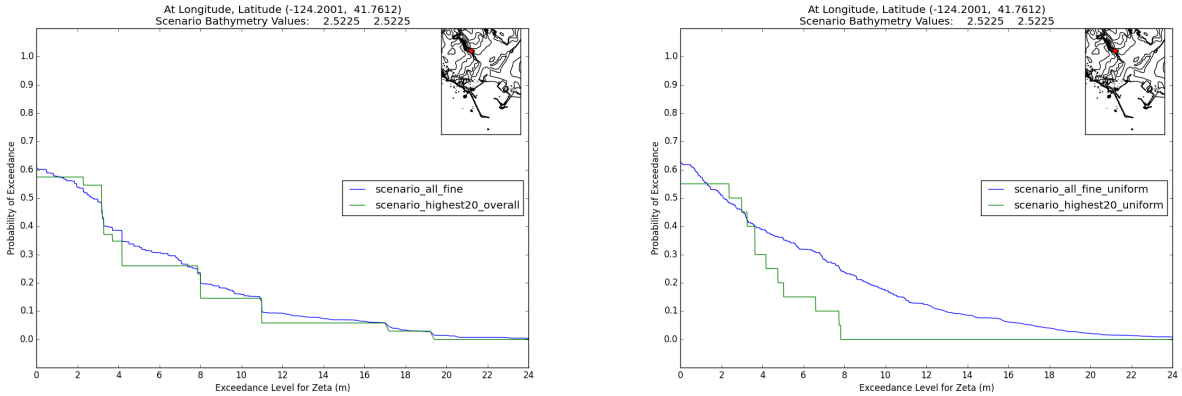

Figure 10: Comparison of the true hazard curves (blue) at one point with approximations obtained using only 20 realizations, chosen as those with the largest probabilities. Left: Assuming the KL probabilities, Right: Assuming uniform probabilities.

We use ζ to represent $h_{\max}$ for points on land (initial bathymetry $B0 \geq 0$) and $h_{\max} + B0$ for points initially in the water ($B0 < 0$). Figure 11 shows the highest 20 strategy compares better to the 400 fine for KL probabilities than it does for uniform probabilities. Since, this highest 20 probability strategy is sensitive to the conditional probabilities used, we can not recommend it in general, though it can be very good when the 20 used represent a large fraction of the conditional probability.

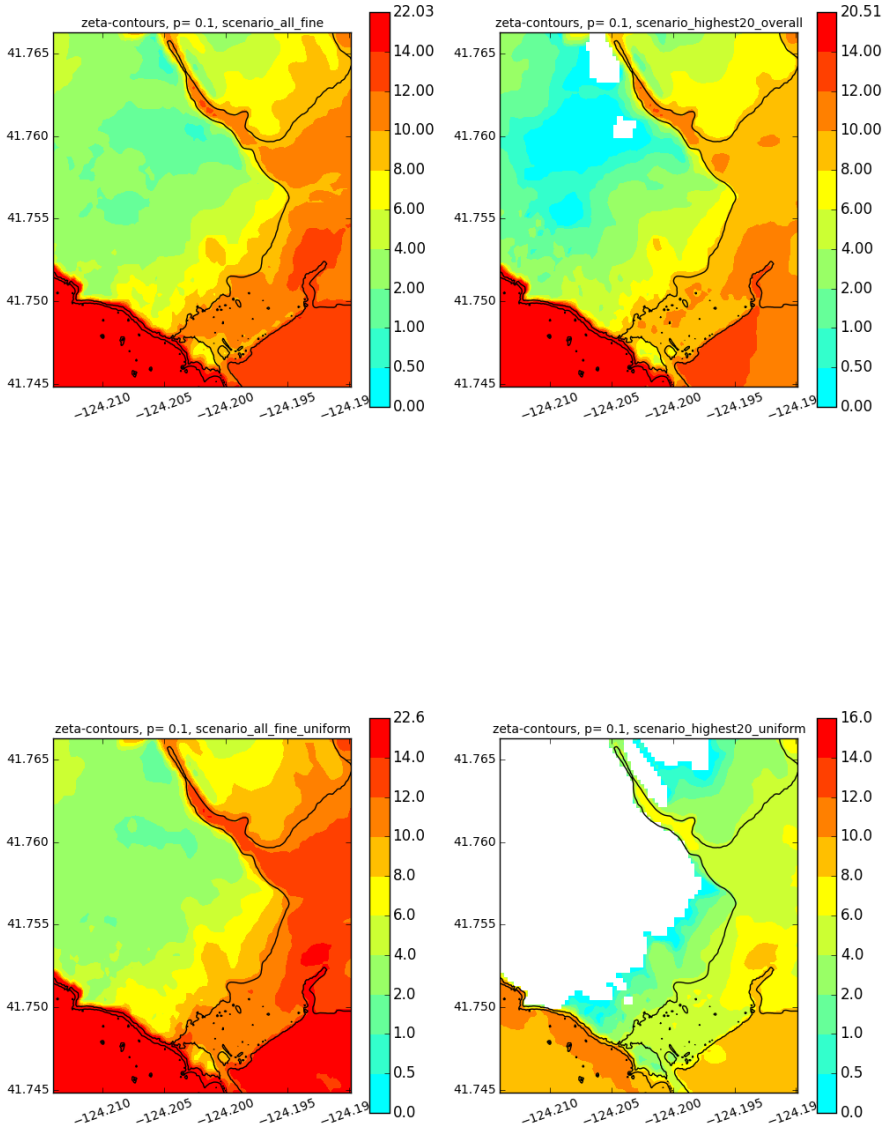

Figure 11: Comparison of the true hazard maps (left) for $p = 0.1$ with those obtained using only 20 realizations (right), chosen as those with the largest probabilities. Top: Assuming the KL probabilities, Bottom: Assuming uniform probabilities.

7.2 All coarse grid strategy

An obvious question is how can information from the 400 coarse grid runs be used to approximate the 400 fine grid results. Without using our budget yet of 20 fine grid runs, we simply compared the 400 coarse to the 400 fine. Results of this can be seen in Figures 12, 13, 14, and 15 by comparing the plots labeled `scenario_all_fine` with those labeled `scenario_all_coarse`.

The conclusion is that the 400 coarse alone can not approximate well at all the 400 fine. However, this comparison gave insight into what has to be done in order to use the 400 coarse to build a strategy. In particular, the $h_{\max}$ value of a coarse grid run represents the amount of water above the *coarse* bathymetry, whereas the $h_{\max}$ value of a fine grid run represents the amount of water above the *fine* bathymetry. Since the two bathymetries can differ at the same physical location, it is imperative that we first "clean" the coarse grid runs so the coarse $h_{\max}$ is modified to represent the amount of water from the coarse grid run that was above the fine grid bathymetry. This is an easy post-processing step, and is the basis of the strategy discussed in Section 7.3.

7.3 Modified coarse grid strategy

This strategy uses all 400 coarse runs that have been modified to measure $h_{\max}$ at each of the points in the 88x78 rectangular grid or on either of the two transects of Figure 1 above the fine grid bathymetry rather than the coarse grid bathymetry. To describe how we make this modification, it is first necessary to introduce some notation.

We let B_f and B_c denote the original bathymetry on the fine and coarse grids at the fine grid point locations, respectively. Recall that B_c is the same length as B_f since we use piecewise constant interpolation to fill in all the fine points within a coarse cell. The difference D_B of these values will be $D_B = B_f - B_c$. We note that if subsidence or uplift of the original bathymetry happened, then B_f and B_c would change by the same amount, and D_B would be unaffected; hence, we can use the original bathymetry for our purposes.

We desire to create a modified coarse grid value for each of the given coarse grid values. Recall, that the coarse grid values have been recorded at all fine grid locations using piecewise constant interpolation. At most locations, we would like the resulting value of eta for the modified value (η_m) to be the same as the η_c of the coarse grid value, so we change only the $h_{\max}$ value and not the surface of the water. We denote the $h_{\max}$ coarse grid and modified coarse values as h_c and h_m, respectively. Then $\eta_c = h_c + B_c$ and $\eta_m = h_m + B_f$. Setting these η values equal means in most cases we should choose the modified coarse grid value as

$$h_m = h_c - D_B\,.$$

There are a few special cases where we can not use the equation above to set h_m. Clearly, if $D_B > h_c$, we can not allow h_m to become negative (the water reaches the coarse bathymetry level, but not the fine level), so h_m is set to 0. Another special case is when $h_c = 0$. 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. Over the region of interest and on land ($B0 \geq 0$), we define η_T as the maximum η_c value where $h_c > 0$; that is, the threshold where it is possible to see flooding. Then we set $\eta_m = \eta_T$ ($h_m = \eta_T - B_f$) or $\eta_m = \eta_c$ ($h_m = B_c - B_f$) depending on whether $\eta_T \leq B_c$ or $\eta_T > B_c$, respectively. With these definitions, h_m will be set as,

$$h_m = \begin{cases} h_c - D_B & \text{if } 0 \leq D_B < h_c \\ 0.0 & \text{if } 0 < h_c \leq D_B \\ h_c - D_B & \text{if } D_B < 0.0 < h_c \\ \eta_T - B_f & \text{if } h_c = 0.0 \text{ and } B_f < \eta_T \leq B_c \\ B_c - B_f & \text{if } h_c = 0.0 \text{ and } B_f < B_c < \eta_T \end{cases}.$$

Two examples are given in Appendix A that show how the modified coarse data compares to both the original coarse data and the fine data. The first example is an 9.2 Mw event, called run 20. The second

is an 8.6 Mw event, called run 5. Both examples verify that the modified coarse data better approximates the fine data than does the original coarse data. This gives hope that the modified coarse grid strategy will approximate the fine grid probability products better than the coarse grid strategy did.

7.4 Pseudo fine grid strategy

This strategy begins by clustering the 400 modified coarse grid runs into 20 non-overlapping clusters. The clusters may contain different numbers of runs. Each cluster will have one run designated as its centroid. We use two different ways of clustering described in Sections 5.1 and 5.2.

Let c_j represent the $h_{\max}$ for a modified coarse grid run that has been assigned to a cluster. Let c and f represent the $h_{\max}$ of the modified coarse grid run and fine grid run, respectively, of the cluster's centroid. Recall that c_j, c, and f are vectors of the length of the fine grid (for example 88x78 for the 2D rectangular fine grid). Then, the pseudo fine grid strategy turns c_j into a pseudo-fine value, p_j, by the simple assignment

$$p_j = c_j + (f - c)\,.$$

Note, if c_j was for the cluster centroid ($c_j = c$), then the pseudo fine value is the same as the fine value ($p_j = f$). So, the 20 centroids become our budget of allowed fine grid runs. These are supplemented by the 380 non-centroid pseudo fine grid runs with $h_{\max}$ values defined by the p_j's.

Two examples are given in Appendix B that show how the pseudo fine data produced with both the etamean-etamax and the dtopo clustering techniques compares to the fine data. The first example is an 8.6 Mw event, called run 5. The second is an 9.2 Mw event, called run 20. Both examples verify that the pseudo fine data from both of the clustering techniques compares fairly well with the fine data. This gives hope that the pseudo fine strategies will also approximate the fine grid probability products fairly well.

7.5 Centroid strategies

We will see in Section 8 that the pseudo fine grid strategy described in Section 7.4 works quite well to approximate the 400 fine grid results when the cluster centroids used to make this strategy were from either of the two clustering techniques. In some cases, the eta mean - eta max clustering was best, and in others the dtopo clustering was best.

Here, we investigate whether just the 20 fine grid runs for the centroids are enough to approximate the 400 fine grid runs. Note that the pseudo fine and fine runs are identical at these 20 centroids. So, the question is whether the remaining 380 pseudo-fine runs are necessary to add to a strategy.

The probabilities used for these 20 centroids are simply the sum of all the probabilities of the runs within the cluster of the centroid. So, either the sum of the KL probabilities within the cluster or the sum of the uniform probabilities within the cluster are assigned to the centroid of the cluster.

In Section 9, we compare two centroid strategies, one based on eta mean - eta max clustering and another based on dtopo clustering to the pseudo fine strategy of Section 7.4.

7.6 SVD strategy

In this strategy we exploit the high similarity between the singular modes of the fine-grid runs and those of the coarse-grid runs. The strategy is inspired by existing multi-level or multi-fidelity strategies; see recent works [22, 8], as well as [12, 24]. Complete details are given in Appendix C.

First, we take the $h_{\max} + B0$ values over the rectangular fgmax grid for all the coarse runs with a fixed magnitude, and reshape the 88×78 grid data to 6864×1 vector data. Then we stack 100 such vectors as columns in a matrix A_c of dimensions 6864×100, then compute its singular value decomposition (SVD),

$$A_c = U_c \Sigma_c V_c^T = U_c Y_c \quad \text{where } Y_c = \Sigma_c V_c^T. \tag{1}$$

If we take the first 4 left singular modes (the first 4 columns of U_c) and look at the corresponding first 4 rows of Y_c, we obtain the *coordinates* of each coarse-grid run in terms of the 4 left singular modes. That is, we obtain a 4-dimensional coordinate for each of the 100 runs, organized in the 4×100 matrix $\bar{Y}_c$.

Using $\bar{Y}_c$, we select the fine-grid runs to perform, by determining which runs have the largest relative component in the 4-coordinate system in the singular directions. Here we select 2 such runs for each of the 4 singular modes. Each of these 8 runs form a cluster for this strategy, in a sense.

Taking these 8 fine-grid $h_{\max} + B0$ values, we again reshape and stack them as column vectors to form the matrix B_f, then compute its SVD,

$$B_f = \bar{U}_f \bar{\Sigma}_f \bar{V}_f^T = \bar{U}_f \bar{Y}_f \quad \text{where } \bar{Y}_f = \bar{\Sigma}_f \bar{V}_f^T. \tag{2}$$

Next, we compute a matrix T that transforms $\bar{Y}_c$ to match the first four rows of $\bar{Y}_f$ as much as possible, by solving a least squares problem.

We denote the first four columns of $\bar{U}_f$ as $\hat{U}_f$. Finally, we multiply $\hat{U}_f$ with $T\bar{Y}_c$ to compute the approximation $\hat{A}_f$ of all the fine runs. That is, we compute

$$\hat{A}_f = \hat{U}_f T \bar{Y}_c. \tag{3}$$

Since we used 8 true fine-grid runs, we then set these particular columns of $\hat{A}_f$ to be those of A_f. From each of these 100 columns of $\hat{A}_f$ we subtract off $B0$ over the onshore points to obtain the 100 ζ vectors.

This approximation is done for each of the magnitudes (8.6, 8.8, 9.0, and 9.2). Since we used 8 fine-grid runs for each magnitude, we have used 32 fine-grid runs total, within the computational budget of 40 runs.

In Section 10, we compare this strategy to the 400 all-fine and a strategy of using only 40 fine grid runs (the 10 highest KL probability runs from each of the magnitude categories), and an all-pseudo strategy that used 40 clusters determined from the modified coarse grid runs.

8 Comparisons of coarse, modified coarse, and pseudo strategies

This section compares the 400 coarse, the 400 modified coarse, and the 400 pseudo-fine strategies with the 400 fine results all using the KL probabilities. Figure 12 is a hazard curve comparison at the location indicated by the red dot. This location is in the valley with steep bathymetry on either side. This comparison is typical of other locations across the region of interest, and shows the pseudo-fine strategy to be a good approximation. For this particular hazard curve the weighted difference in ζ values (as define in Section 6) was 1.37m for scenario-all-coarse, 0.36m for scenario-all-pseudo, 1.62m for scenario-all-coarse-mod, and 0.37m for scenario-all-pseudo-dtopo.

The zeta (ζ) contours in Figure 13, the differences of these ζ contours in Figure 14, the weighted differences of these ζ contours in Figure 15, as well as the probability contours of ζ exceeding 2 meters in Figure 16 and the differences of these probability contours in Figure 17 show that the modified coarse is generally an improvement over the coarse strategy, but still not sufficient to approximate the 400 fine results. These figures, like the hazard curve in Figure 12, show the pseudo strategies using the KL probabilities are quite good. In Section 9, we also consider the pseudo strategies with both the KL and uniform probabilities and compare with the centroid strategies.

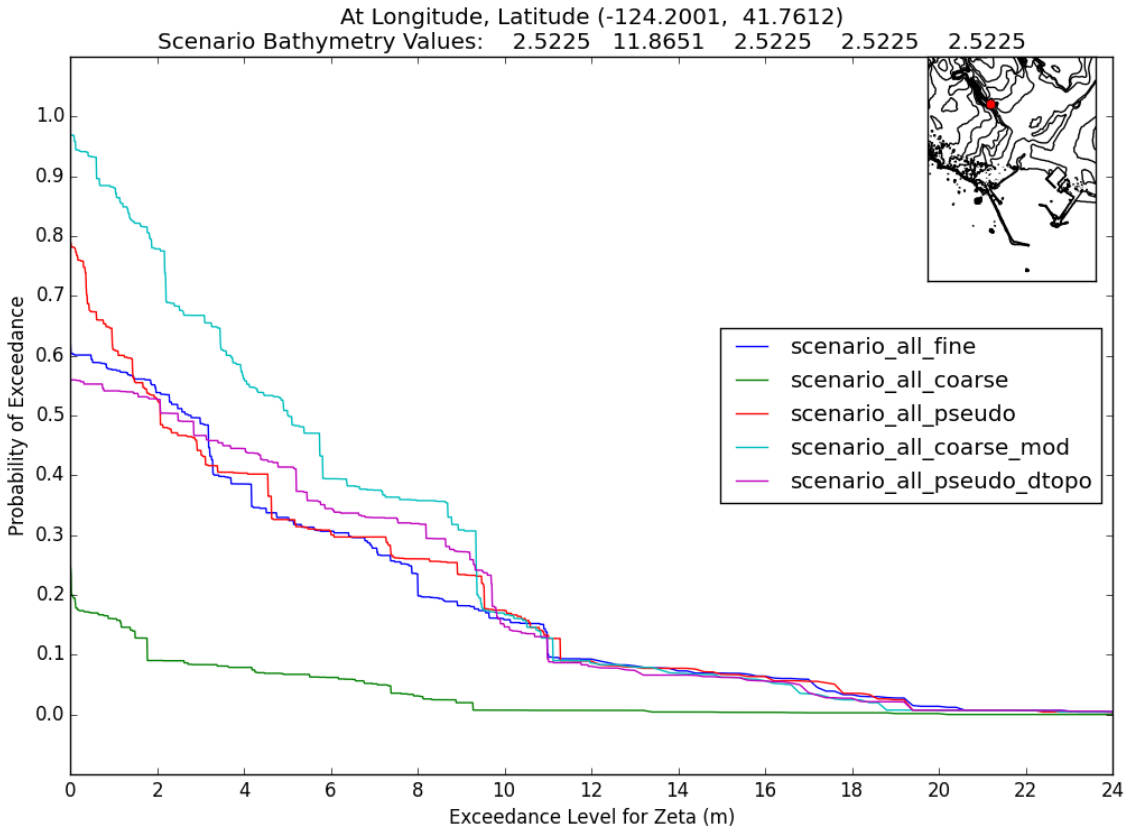

Figure 12: Hazard curves comparison of coarse and pseudo strategies using KL probabilities. The dark blue curve is the "true" hazard curve, as computed from the 400 fine grid simulations.

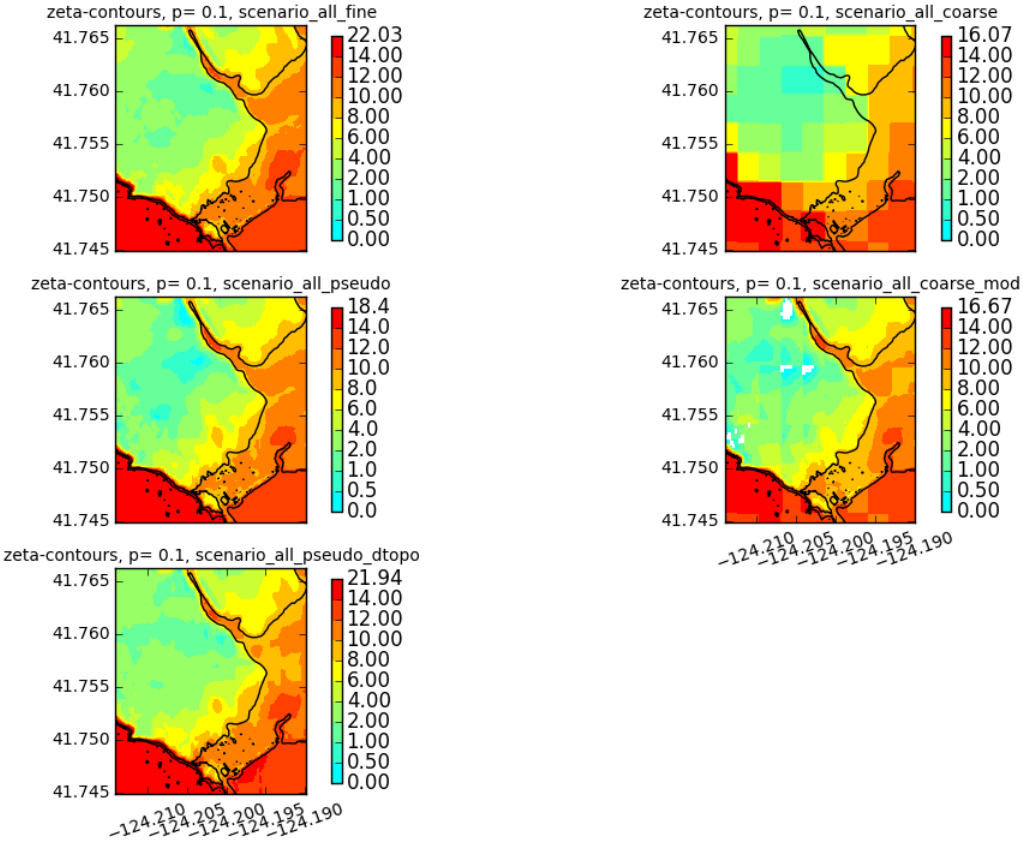

Figure 13: Zeta-contours comparison of coarse and pseudo strategies using KL probabilities

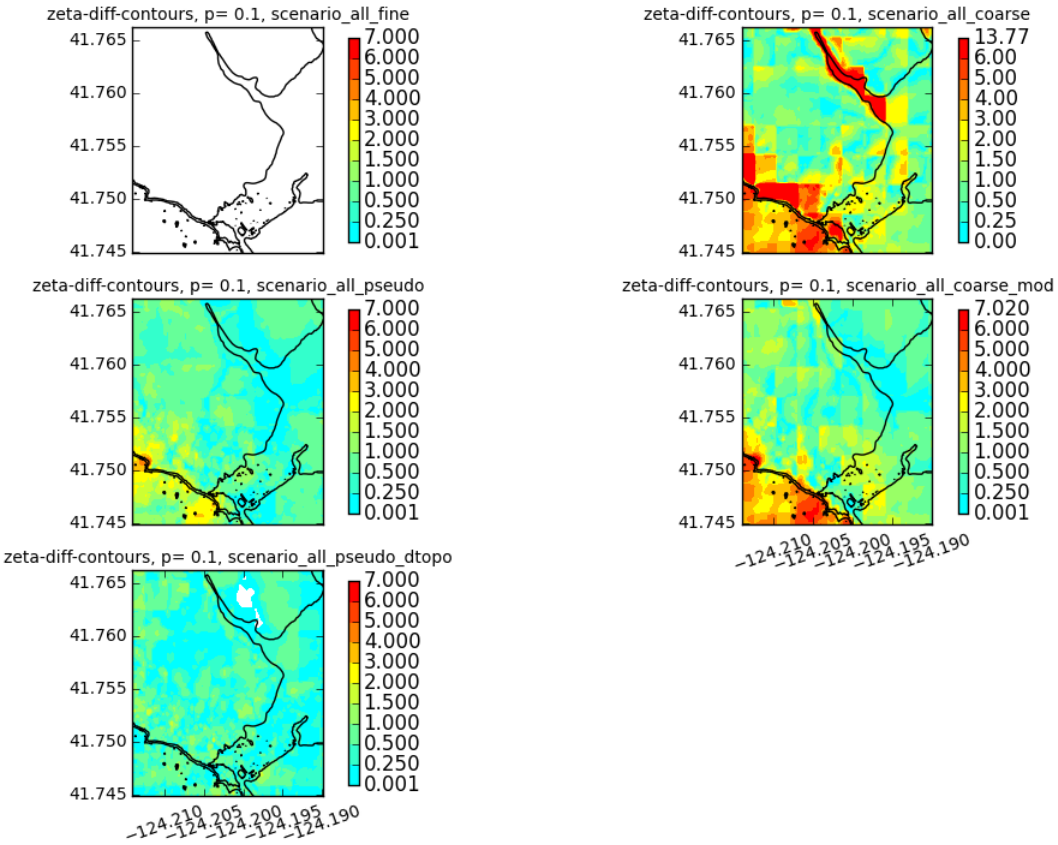

Figure 14: Zeta difference contours of coarse and pseudo strategies using KL probabilities

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 15 were (5.03,0.62) for scenario-all-coarse, (1.24,0.16) for scenario-all-pseudo, (2.46,0.39) for scenario-all-coarse-mod, and (1.71,0.24) for scenario-all-pseudo-dtopo. Notice also the colorbar has changed slightly from Figure 14 to reflect these lower values.

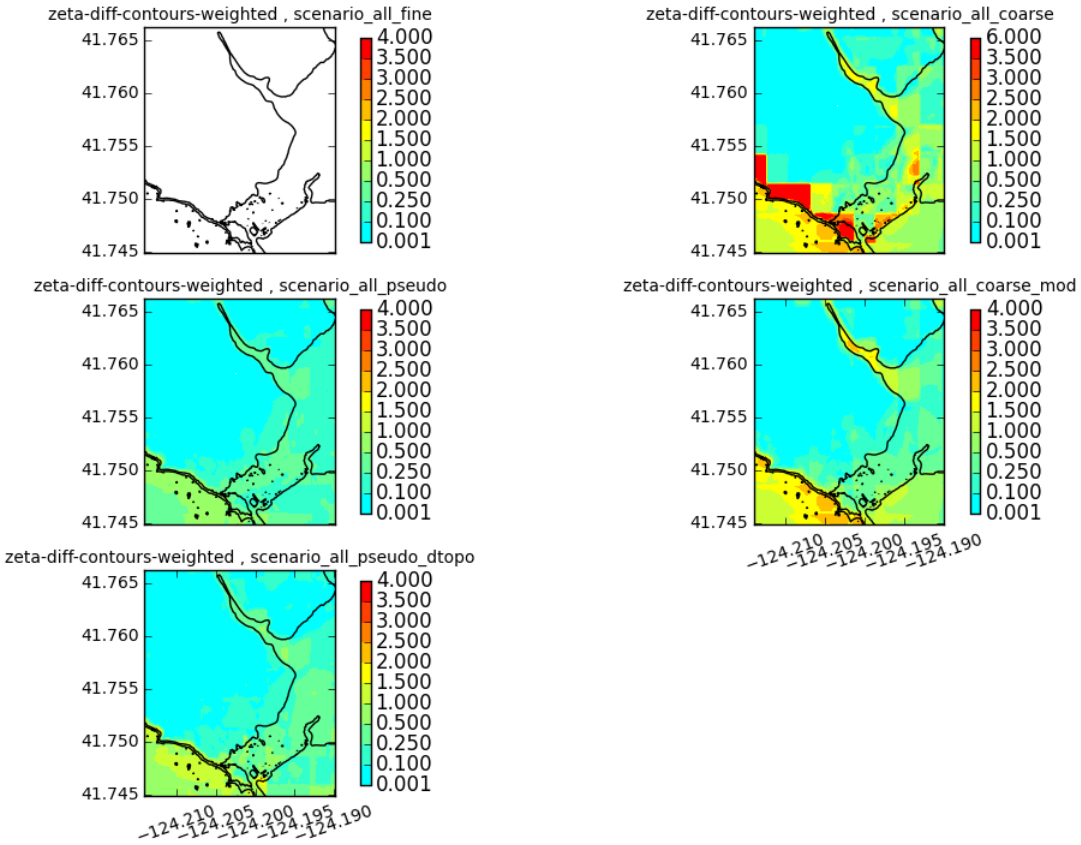

Figure 15: Weighted zeta difference contours of coarse and pseudo strategies using KL probabilities

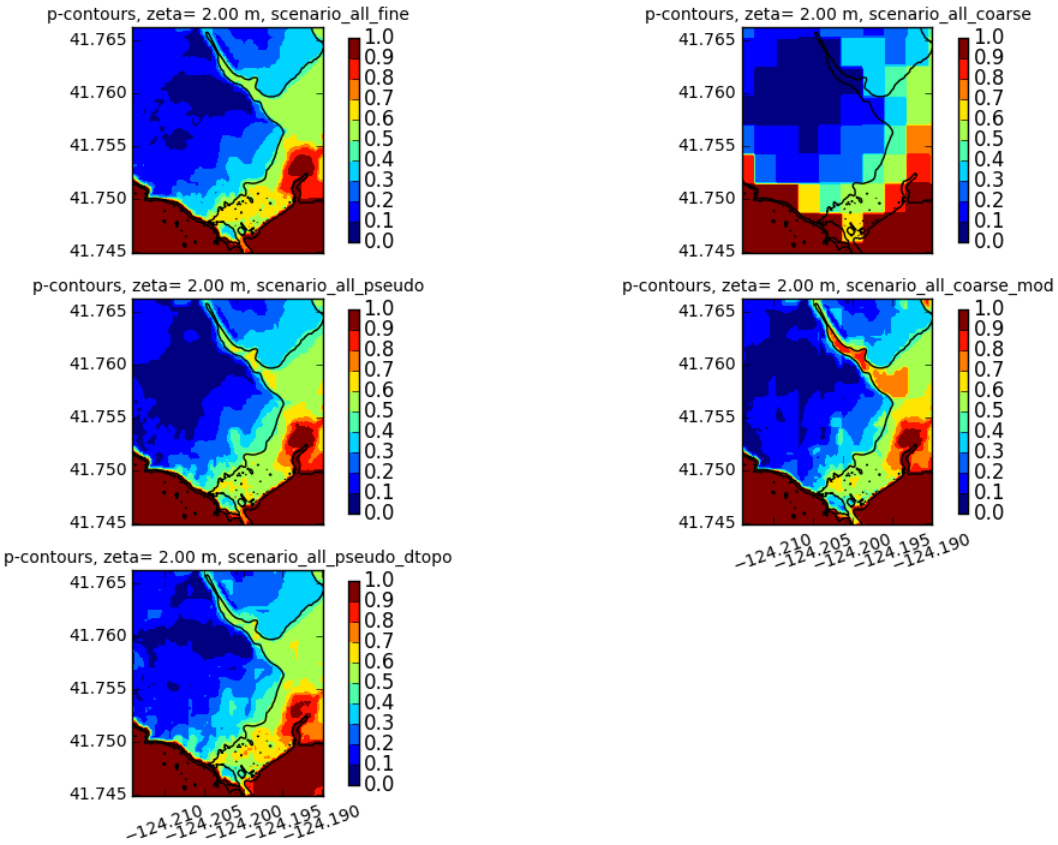

Figure 16: p-contours (2m.) for coarse, modified coarse, pseudo strategies using KL probabilities

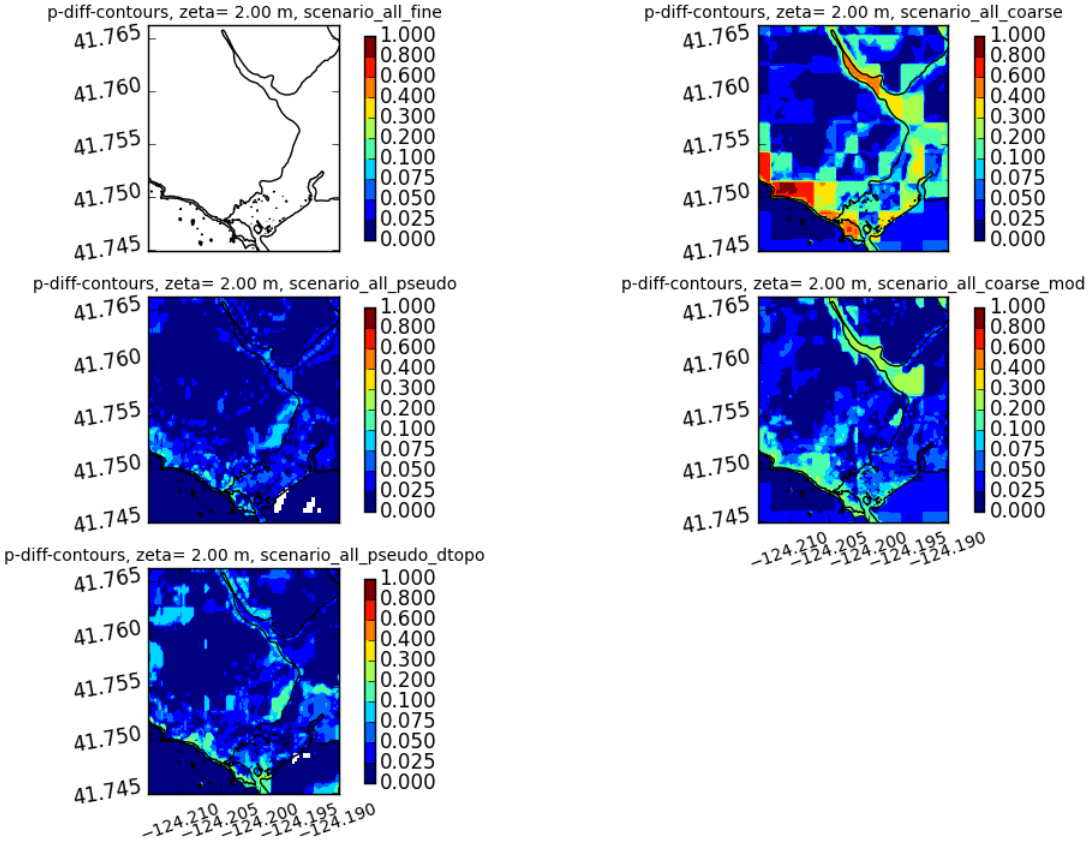

Figure 17: p-contours (2m.) differences for coarse, modified coarse, pseudo strategies using KL probabilities

9 Comparisons of centroid and pseudo strategies

Figure 18 compares the 20-centroid and pseudo strategies. The top plot in the figure is for the KL probabilities and the bottom plot is for the uniform probabilities. For the top hazard curve, the weighted difference in ζ values was 0.36m for scenario-all-pseudo, 0.34 for scenario-dzCC-dtopo-etamax-20cls, 0.37 for scenario-coarse-eta-20cls, and 0.37m for scenario-all-pseudo-dtopo. For the bottom hazard curve, the weighted difference in ζ values was 0.18m for scenario-all-pseudo-uniform, 0.53 for scenario-dzCC-dtopo-20cls-uniform, 0.30 for scenario-c-eta-20cls-uniform, and 0.22m for scenario-all-pseudo-dtopo-uniform. The 20-centroid plots clearly show the discrete steps in the hazard curve due to only 20 events. This is the nature of hazard curves – these steps only smooth out as more events are added. For this reason alone, the pseudo strategies which use all 400 events are more representative of the hazard.

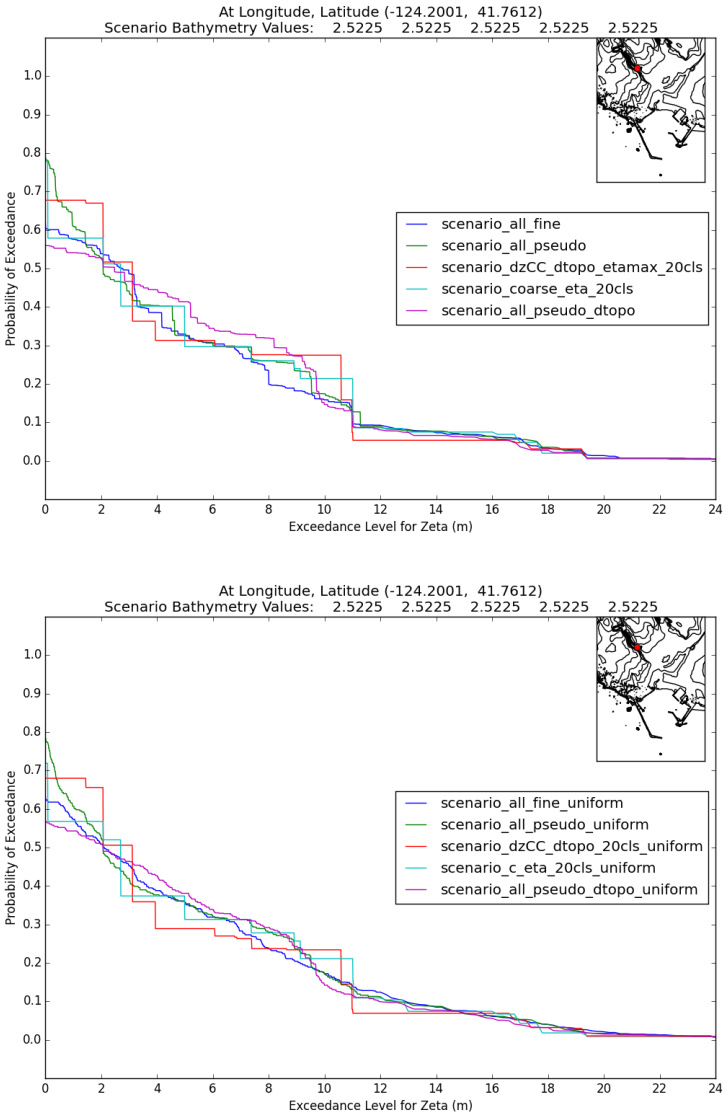

Figure 18: Hazard curves comparison of centroid and pseudo strategies. Top: KL probabilities, Bottom: uniform probabilities. The dark blue curve is the "true" hazard curve, as computed from the 400 fine grid simulations.

Figures 19, 20, and 21 compare the zeta (ζ) contours for $p = 0.1$, their differences, and their weighted differences to the 400 fine for the 20 centroid and pseudo strategies when the KL probabilities are used. Plots in the first column of Figure 20 clearly show that adding in all the 380 pseudo-fine runs to the 20 centroids improves the approximation dramatically when dtopo clustering is used. The plots in the second column also shows some improvement when the 380 are added when eta mean - eta max clustering is used.

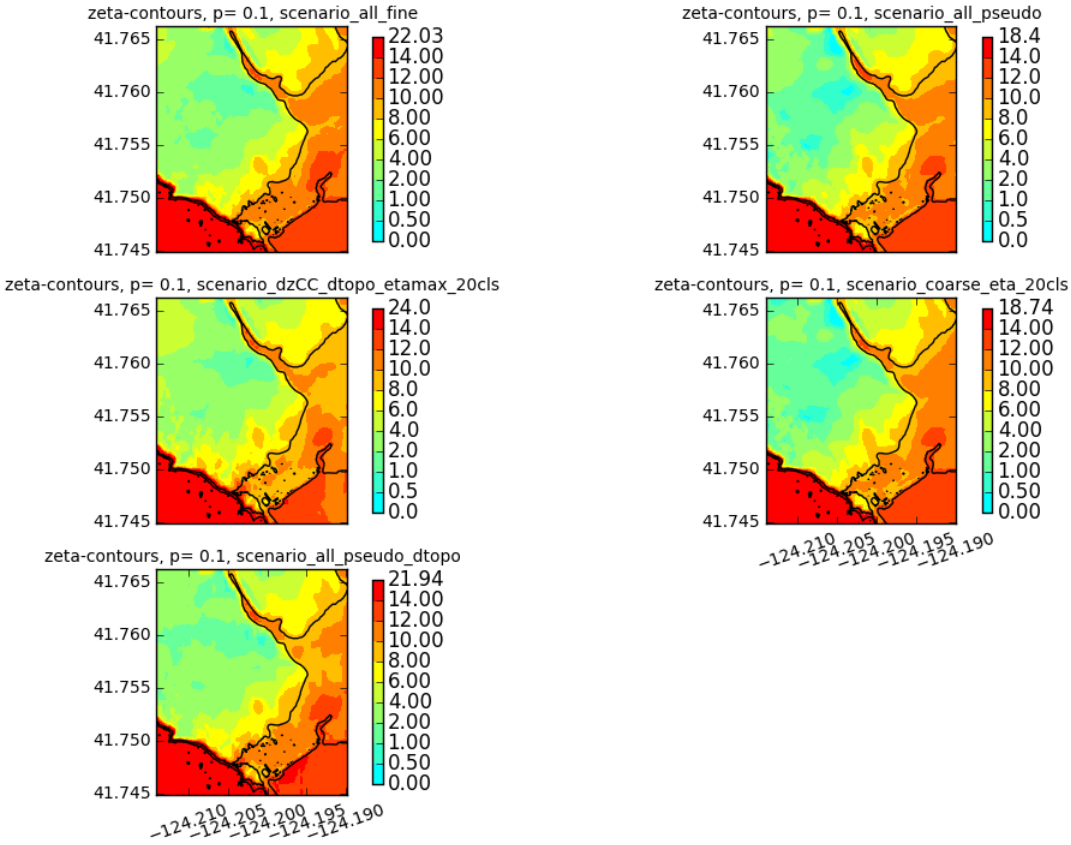

Figure 19: Zeta-contours comparison of centroid and pseudo strategies using KL probabilities

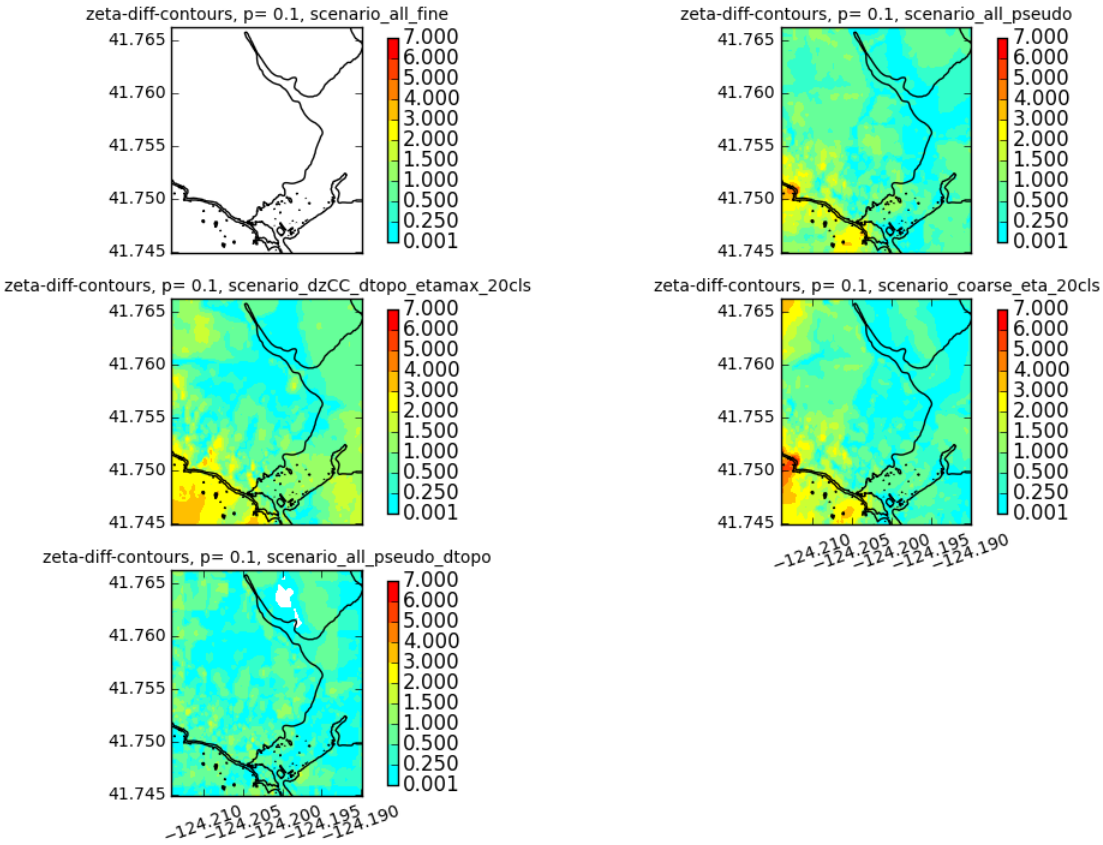

Figure 20: Zeta difference contours of centroid and pseudo strategies using KL probabilities

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 21 were (1.24,0.16) for scenario-all-pseudo, (2.43,0.46) for scenario-dzCC-dtopo-etamax-20cls, (1.34,0.27) for scenario-coarse-eta-20cls, and (1.71,0.24) for scenario-all-pseudo-dtopo. Notice also the colorbar has changed slightly from Figure 20 to reflect these lower values.

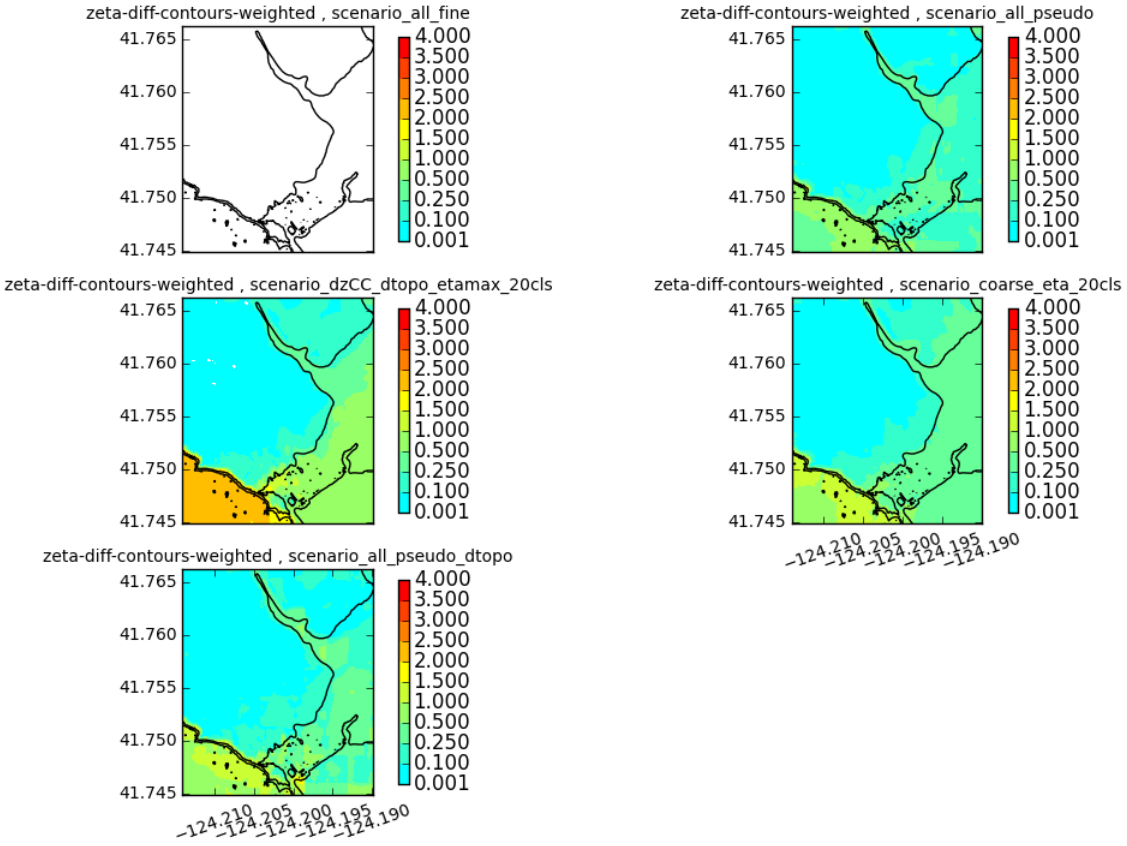

Figure 21: Weighted zeta difference contours of centroid and pseudo strategies using KL probabilities

Figures 22, 23, and 24 compare the zeta (ζ) contours for $p = 0.1$, their differences, and their weighted differences to the 400 fine for the 20 centroid and pseudo strategies when the uniform probabilities are used. Plots in the first column of Figure 23 clearly show that adding in all the 380 pseudo-fine runs to the 20 centroids improves the approximation dramatically when dtopo clustering is used. The plots in the second column also shows some improvement when the 380 are added when eta mean - eta max clustering is used.

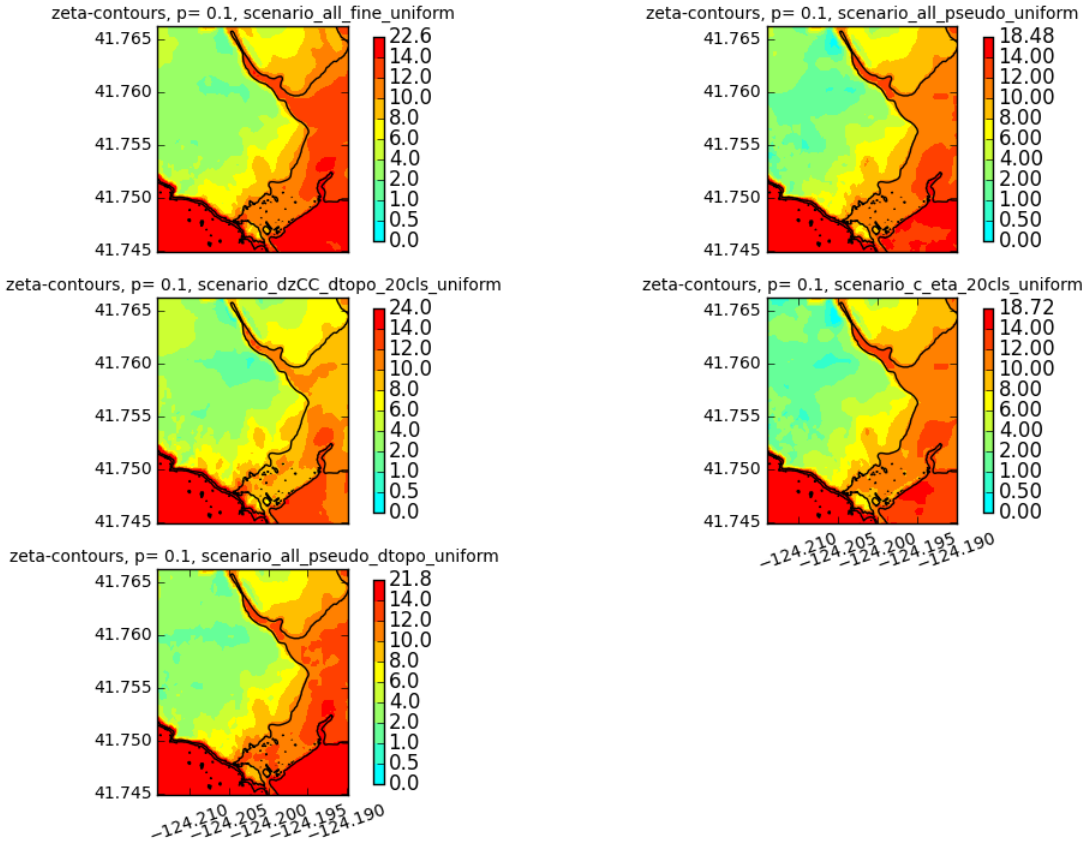

Figure 22: Zeta-contours comparison of centroid and pseudo strategies using uniform probabilities

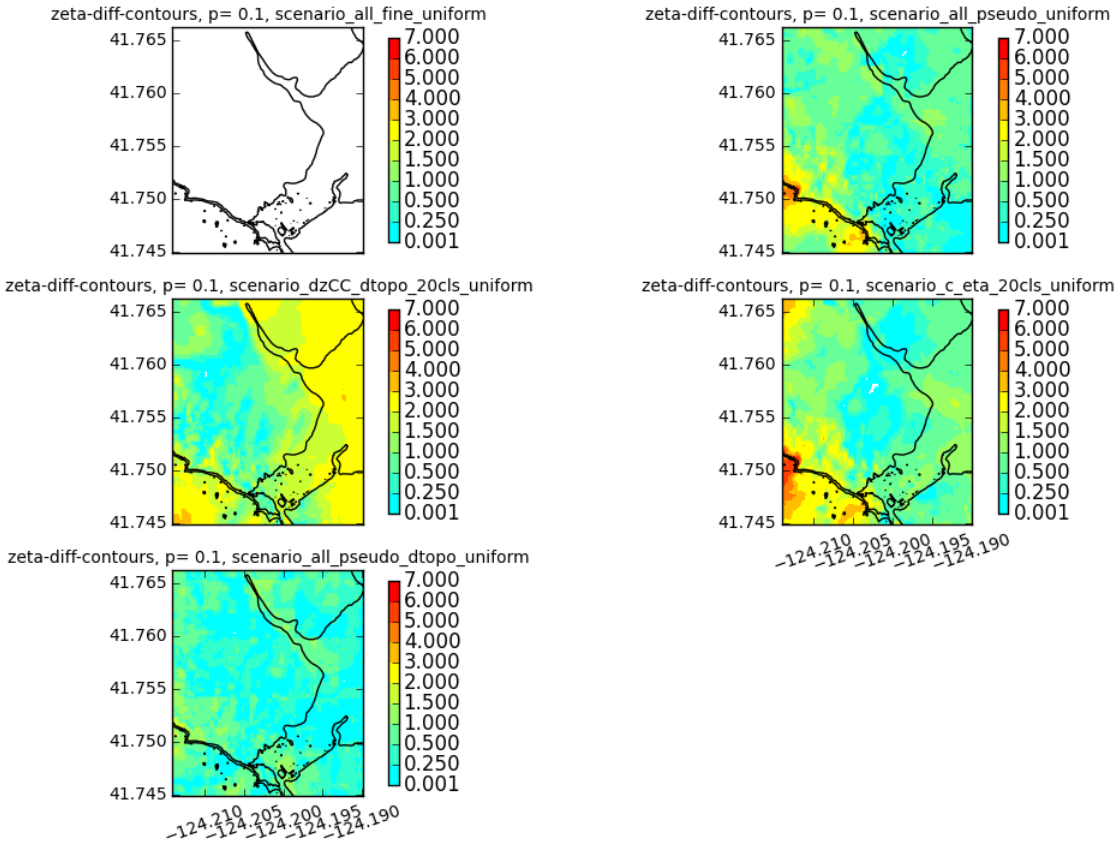

Figure 23: Zeta difference contours of centroid and pseudo strategies using uniform probabilities

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 24 were (0.90,0.12) for scenario-all-pseudo-uniform, (2.83,0.55) for scenario-dzCC-dtopo-20cls-uniform, (0.91,0.21) for scenario-c-eta-20cls-uniform, and (1.71,0.23) for scenario-all-pseudo-dtopo-uniform. Notice also the colorbar has changed slightly from Figure 23 to reflect these lower values.

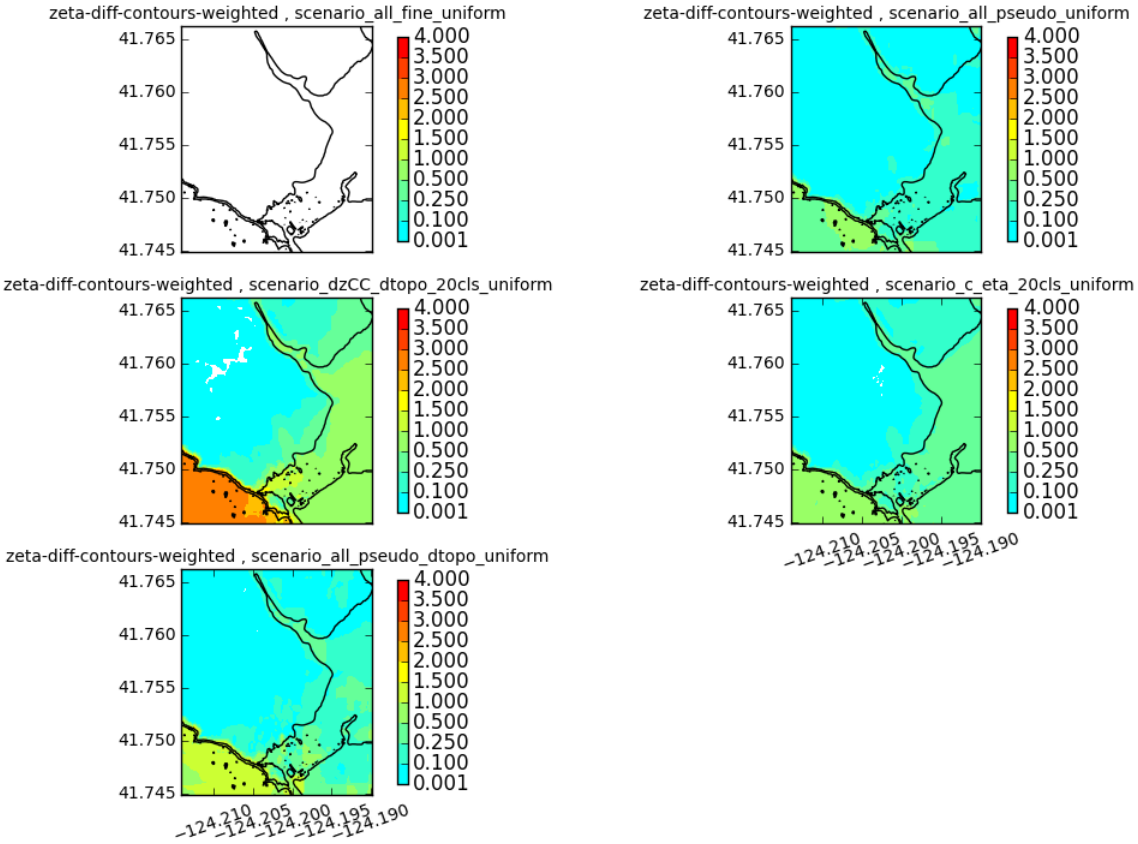

Figure 24: Zeta difference contours of centroid and pseudo strategies using uniform probabilities

Figures 25 and 26 compare the p-contours for exceeding $\zeta = 2$ meters and their differences to the 400 fine for the 20 centroid and pseudo strategies when the KL probabilities are used. Plots in the first column of Figure 26 clearly show that adding in all the 380 pseudo-fine runs to the 20 centroids improves the approximation dramatically when dtopo clustering is used. The plots in the second column also show an overall improvement when the 380 are added and eta mean - eta max clustering is used.

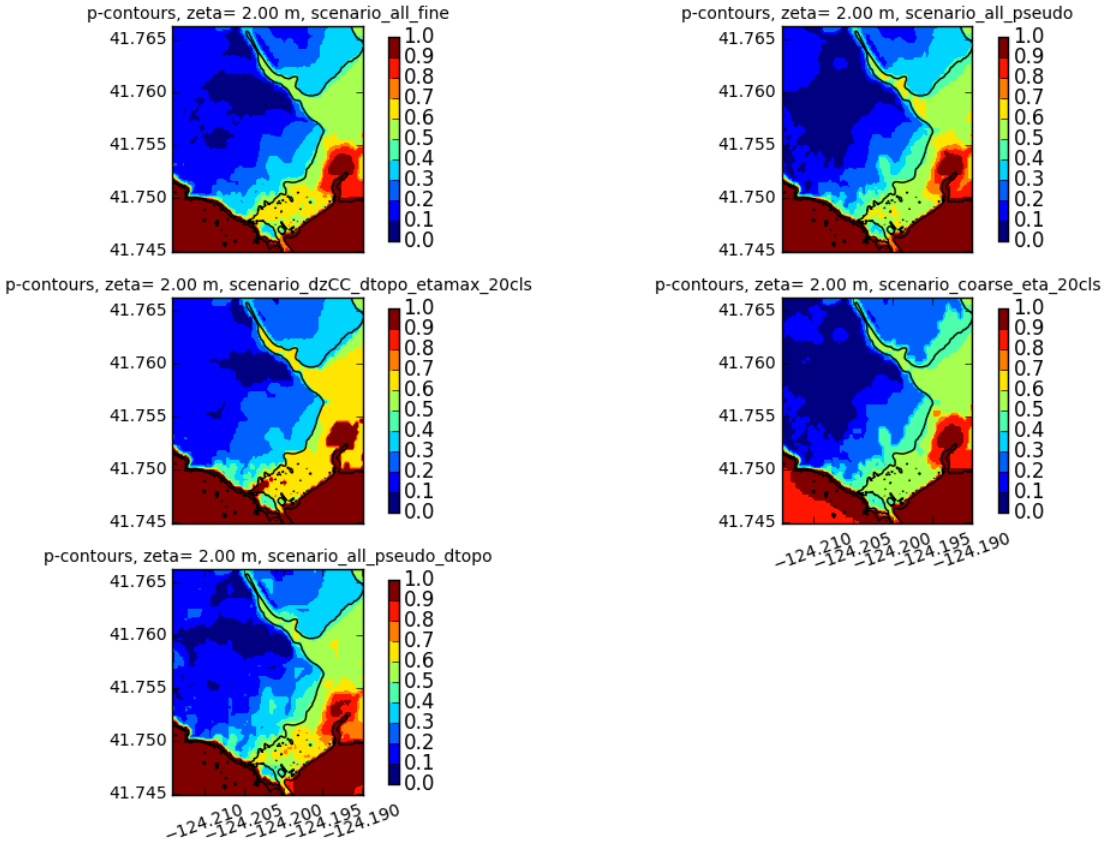

Figure 25: p-contours (2m.) for centroid and pseudo strategies using KL probabilities

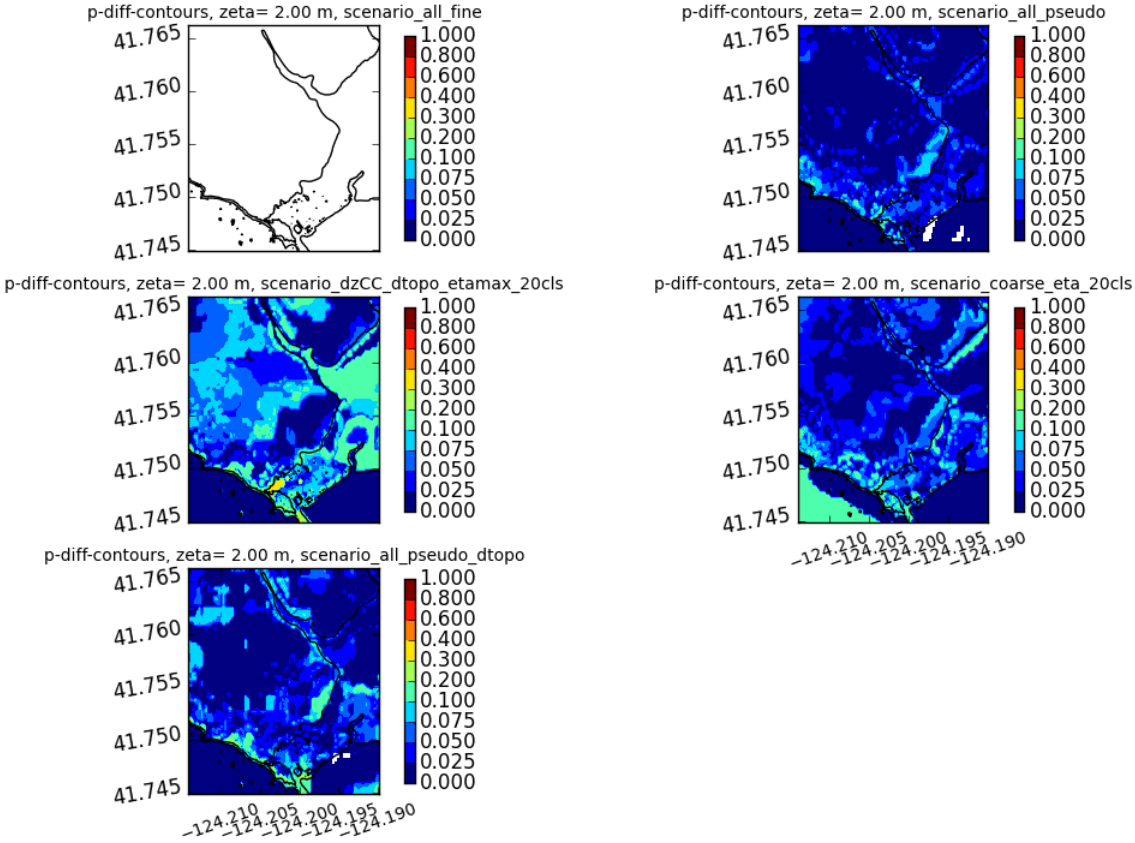

Figure 26: p-contours (2m.) differences for centroid and pseudo strategies using KL probabilities

Figures 27 and 28 compare the p-contours for exceeding $\zeta = 2$ meters and their differences to the 400 fine for the 20 centroid and pseudo strategies when the uniform probabilities are used. Plots in the first column of Figure 28 show that adding in all the 380 pseudo-fine runs to the 20 centroids improves the approximation dramatically when dtopo clustering is used. The plots in the second column also show a dramatic improvement when the 380 are added and eta mean - eta max clustering is used.

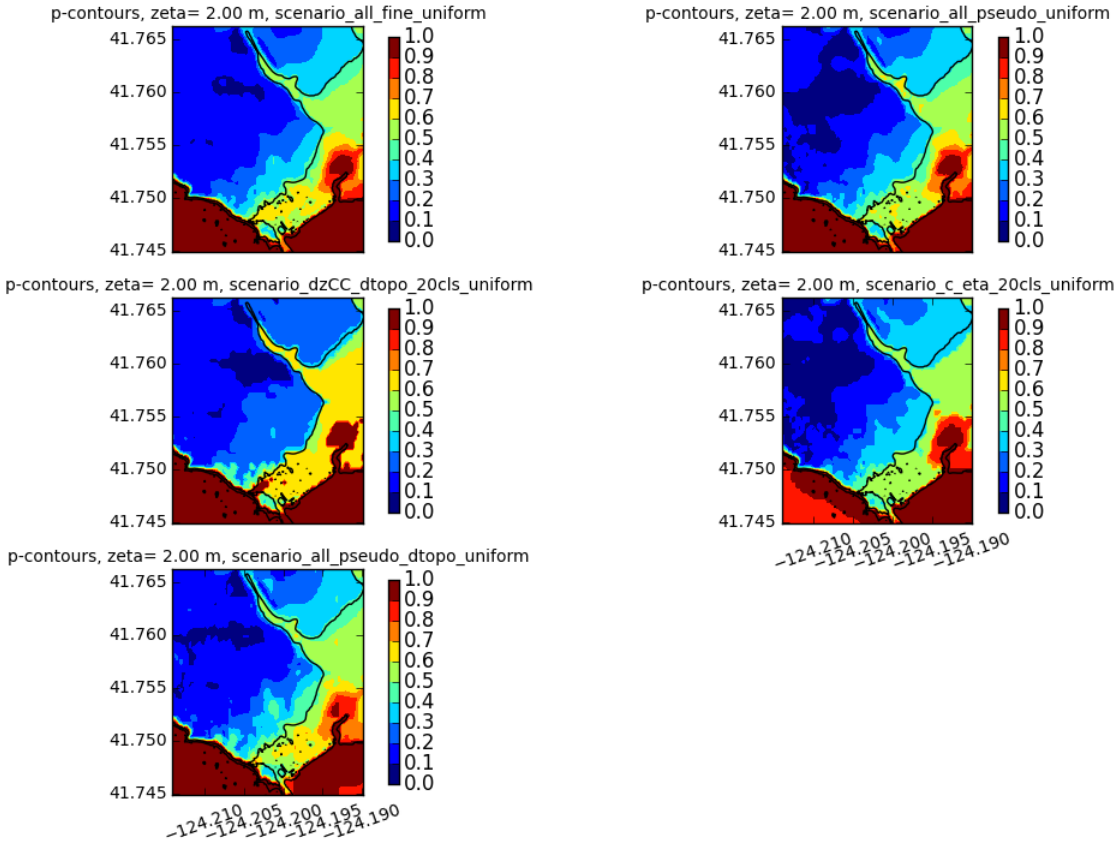

Figure 27: p-contours (2m.) for centroid and pseudo strategies using uniform probabilities

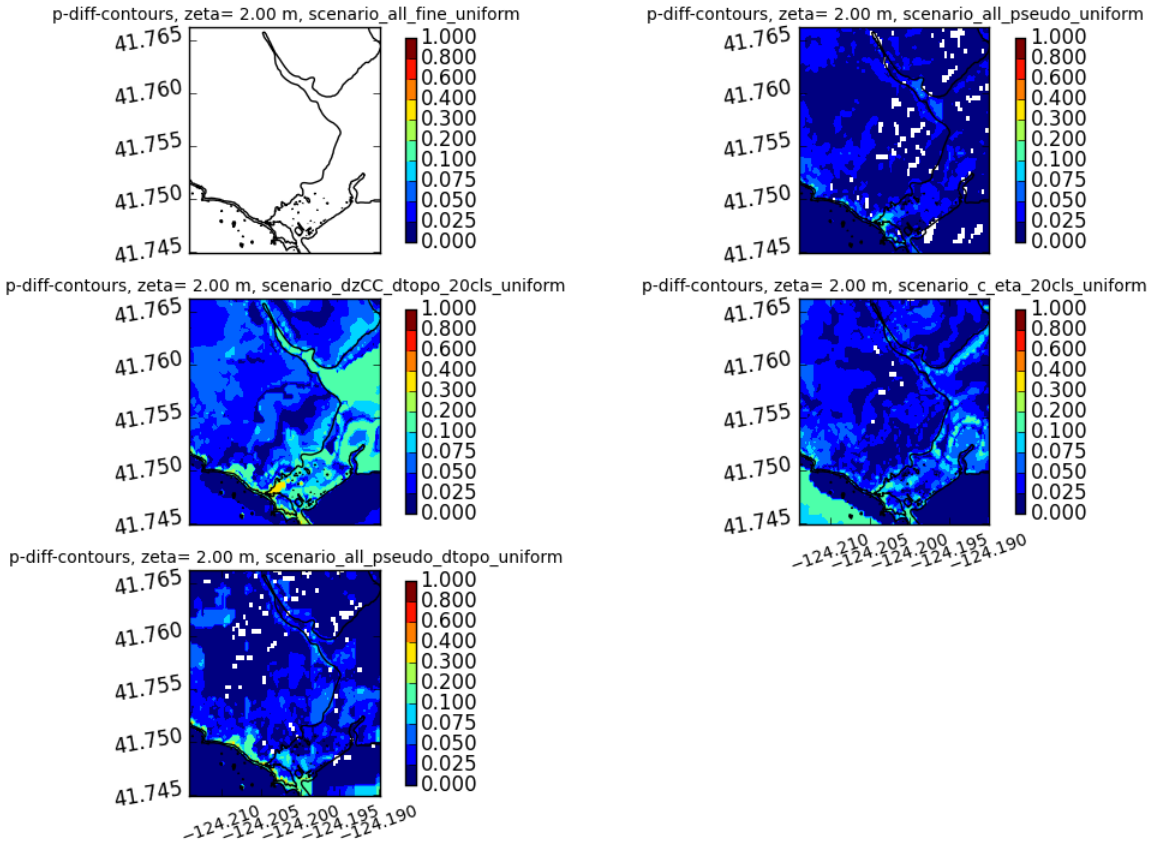

Figure 28: p-contours (2m.) differences for centroid and pseudo strategies using uniform probabilities

In this section, we have shown how the centroid and pseudo strategies compare by selecting a few probability products: a hazard curve at one location, the zeta-contour for $p = 0.1$, a weighted contour of ζ differences across specific probabilities ($p = 0.1$ to $p = 0.9$ in increments of 0.1, and $p = 0.95$), and the p-contour for exceeding $\zeta = 2.0$ meters. In Appendix D, we give the mean and maximum differences of the centroid and pseudo strategies to the fine grid results for the p-contours for exceeding zeta=0, 1, 2, 3, 4, 5, 6, 8, 10, and 12 meters, and for the zeta-contours for p=0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 0.95 when the KL probabilities are used. Another strategy called highest10-each was also compared to the fine grid results for the KL probabilities. This strategy used a budget of 40 fine grid runs (rather than 20) and these were chosen as the 10 highest probabilities from each of the 4 Mw categories (Mw=8.6, 8.8, 9.0, and 9.2). This strategy is included to show as our budget of fine grid runs increases, we can expect the results to improve. In Appendix E, we give the weighted ζ differences of the centroid and pseudo strategies to the fine grid results for the KL probabilities. For comparison, we also include the coarse-mod strategy and some pseudo and centroid strategies that use 40 clusters and give the ζ difference plots for all these strategies. Indeed, using 40 clusters improves the results seen for 20 clusters. In Appendix F, we give the mean and maximum differences for the centroid and pseudo strategies when the uniform probabilities and 20 clusters are used. Transect plot and contour plots for all these cases comparing the centroid and pseudo strategies can be found at http://depts.washington.edu/ptha/FEMA/.

10 Comparisons of highest-10-each, all-pseudoCM-40 and all-SVD-40 strategies

Figure 29 compares the hazard curves for the highest-10-each, all-pseudoCM-40, and all-SVD-40 strategies when the KL probabilities are used. The all-pseudoCM-40 is a pseudo strategy that did the clustering based on the modified coarse runs and used 40 clusters. The all-SVD-40 strategy used 32 fine grid runs (8 from each magnitude category), and since 32 is closer to 40 than to 20, we compare it to our best strategies that use a budget of 40 fine grid runs. The weighted difference in ζ values was 0.22m for scenario-highest-10-each, 0.24m for scenario-all-pseudoCM-40, and 0.56m for scenario-all-SVD-40.

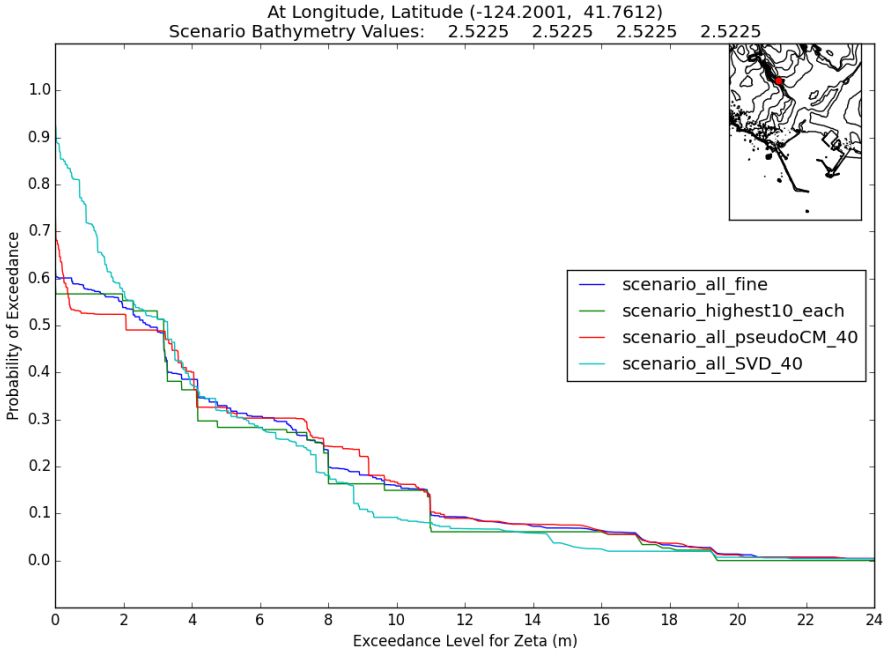

Figure 29: Hazard curves comparison of highest-10-each, all-pseudoCM-40, and all-SVD-40 strategies using KL probabilites

Figures 30, 31, and 32 compare the zeta (ζ) contours for $p = 0.1$, their differences, and their weighted differences to the 400 fine for the highest-10-each, all-pseudoCM-40 and all-SVD-40 strategies when the KL probabilities are used.

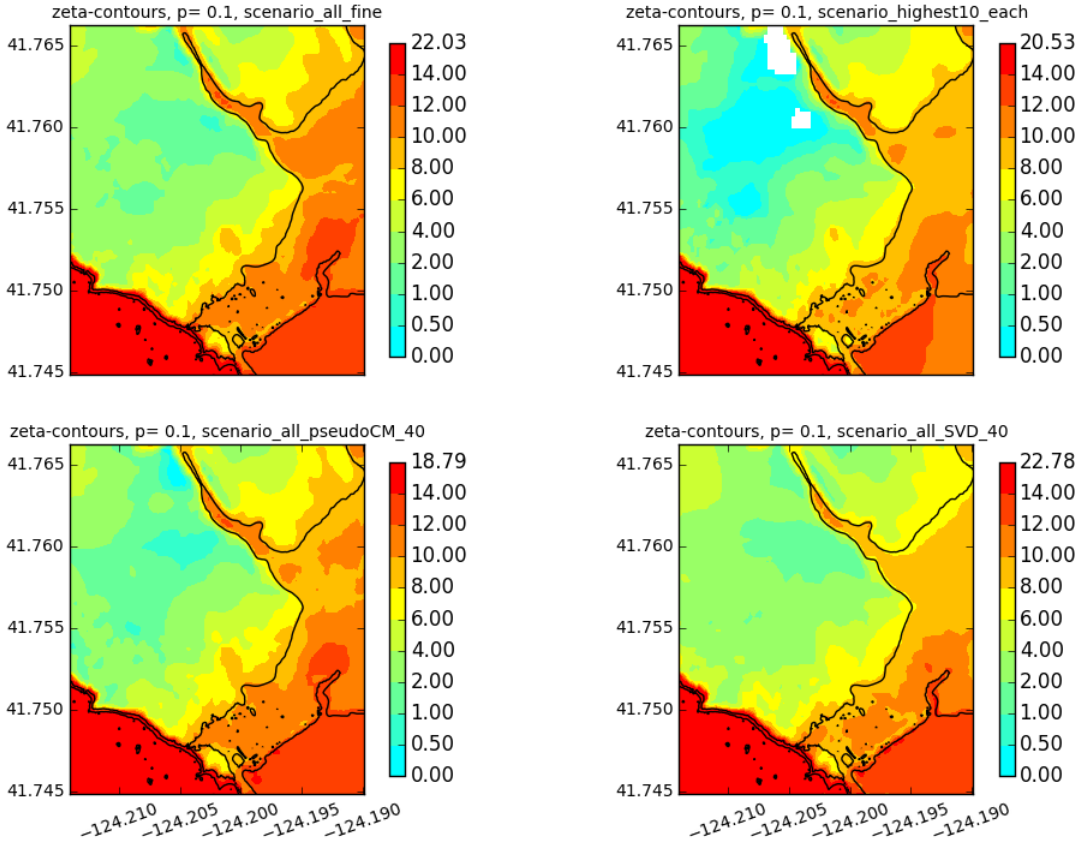

Figure 30: Zeta-contours comparison of the highest-10-each, all-pseudoCM-40 and all-SVD-40 strategies using KL probabilities

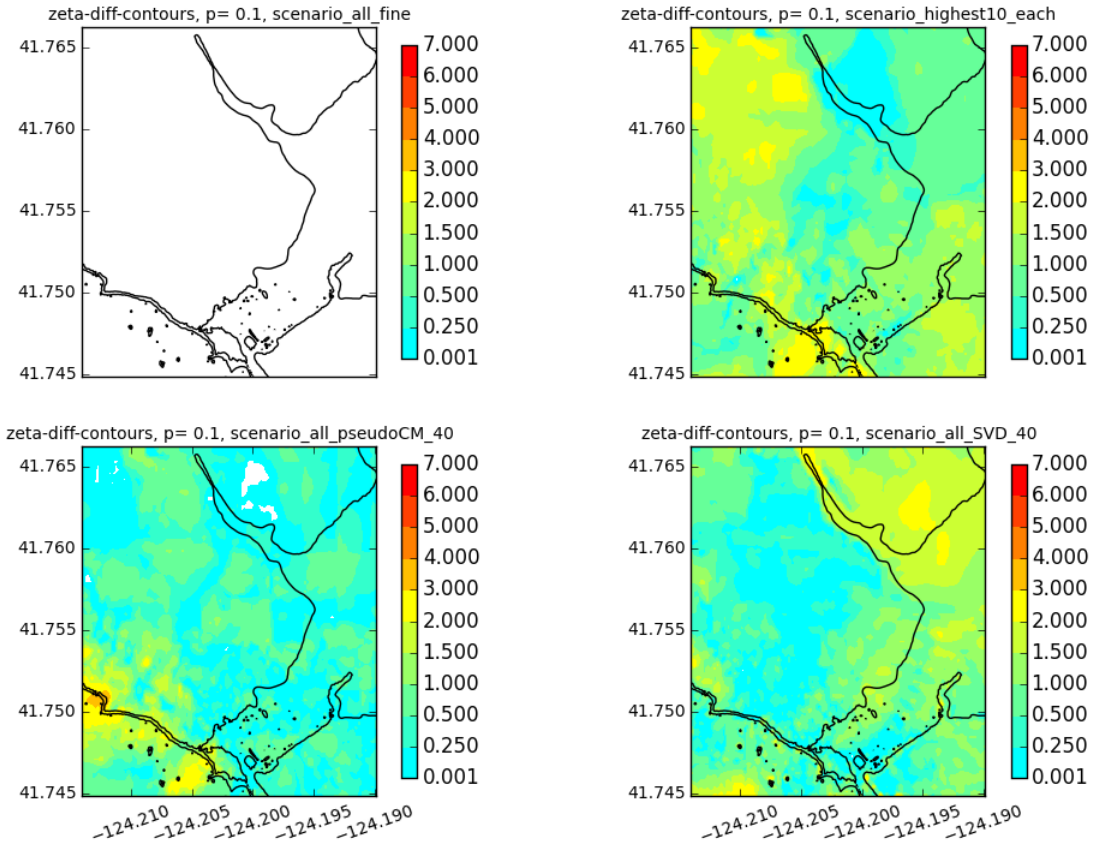

Figure 31: Zeta difference contours of the highest-10-each, all-pseudoCM-40 and all-SVD-40 strategies using KL probabilities

The (maximum, mean) of the weighted ζ differences across the fixed grid in Figure 32 were (0.38,0.11) for scenario-highest-10-each, (0.65,0.14) for scenario-all-pseudoCM-40, and (2.49,0.23) for scenario-all-SVD-40. Notice also the colorbar has changed slightly from Figure 31 to reflect these lower values.

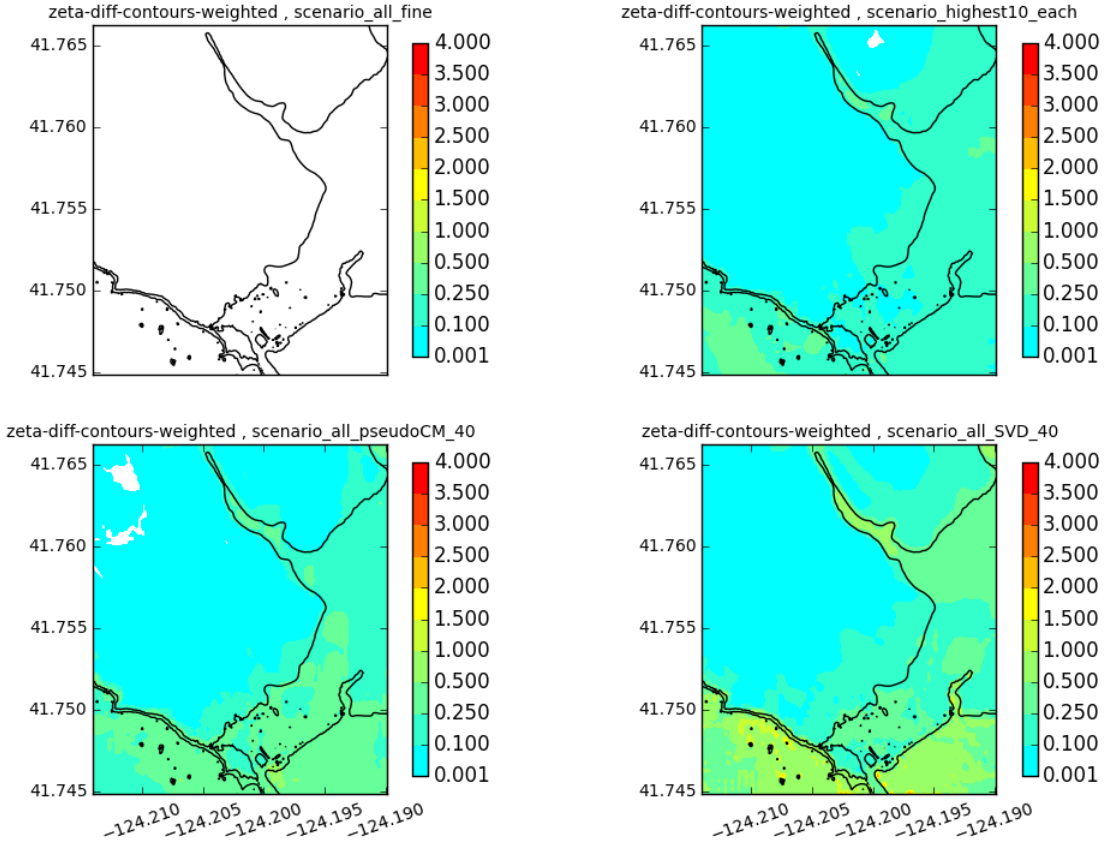

Figure 32: Weighted zeta difference contours of the highest-10-each, all-pseudoCM-40 and all-SVD-40 strategies using KL probabilities

Figures 33 and 34 compare the p-contours for exceeding $\zeta = 2$ meters and their differences to the 400 fine for the highest-10-each, all-pseudoCM-40 and the all-SVD-40 strategies when the KL probabilities are used.

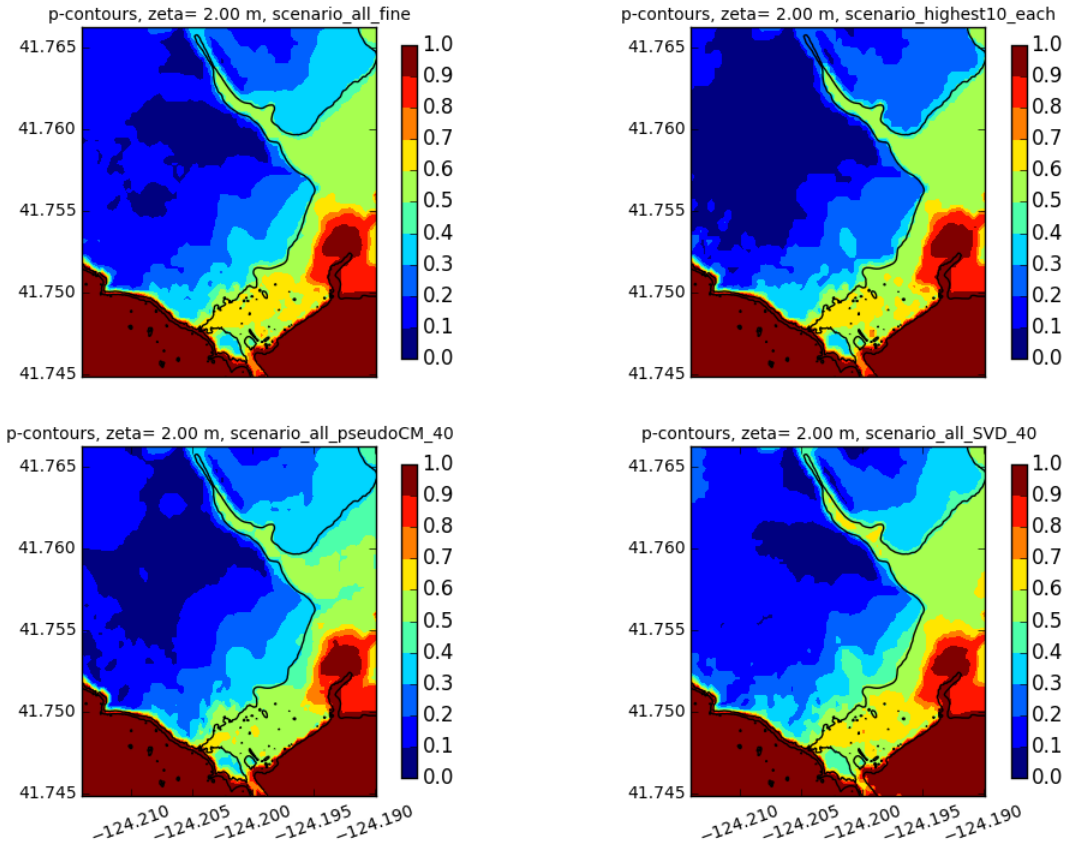

Figure 33: p-contours (2m.) for the highest-10-each, all-pseudoCM-40 and the all-SVD-40 strategies using KL probabilities

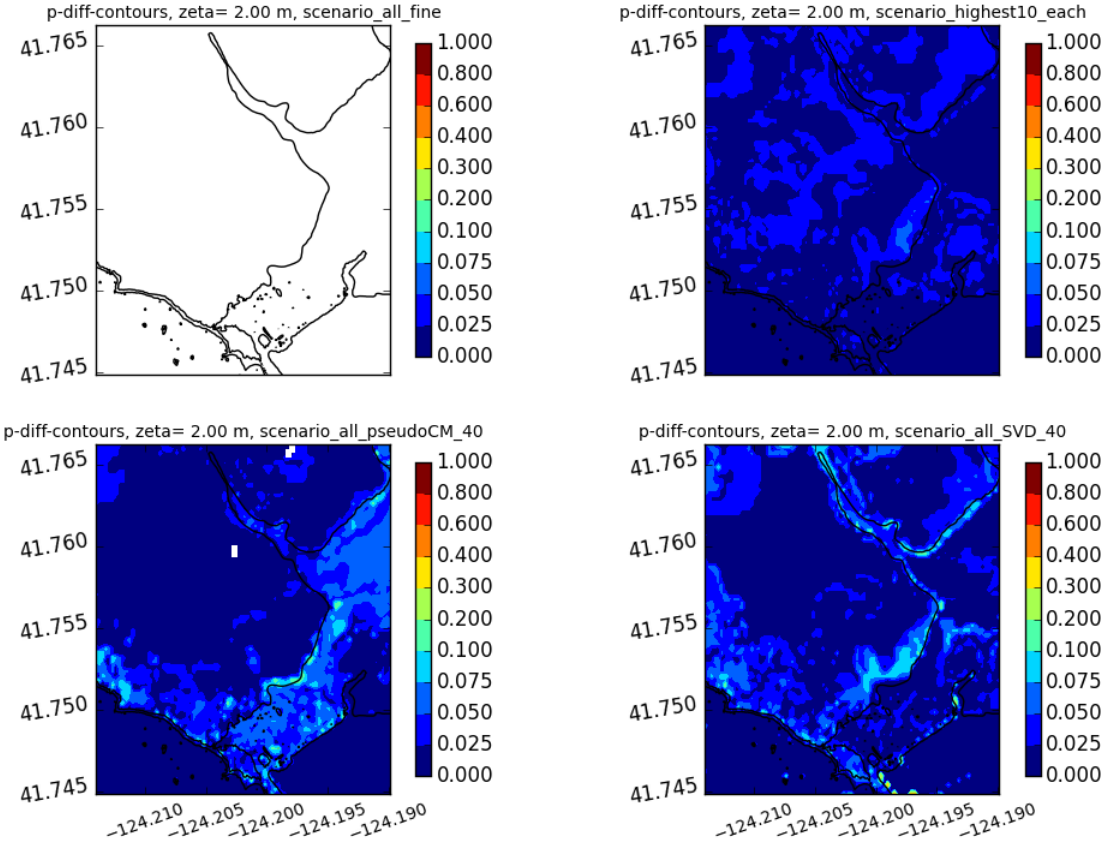

Figure 34: p-contours (2m.) differences for the highest-10-each, all-pseudoCM-40 and the all-SVD-40 strategies using KL probabilities

In this section, we have shown how the highest-10-each, all-pseudoCM-40 and all-SVD-40 strategies compare to all 400 fine when the KL probabilities are used by comparing a few probability products: a hazard curve at one location, the zeta-contour for $p = 0.1$, a weighted contour of ζ differences across specific probabilities ($p = 0.1$ to $p = 0.9$ in increments of 0.1, and $p = 0.95$), and the p-contour for exceeding $\zeta = 2.0$ meters. In Appendix D, we give the mean and maximum differences of these strategies to the fine grid results for the p-contours for exceeding zeta=0, 1, 2, 3, 4, 5, 6, 8, 10, and 12 meters, and for the zeta-contours for p=0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 0.95 when the KL probabilities are used. In Appendix E, we give the weighted ζ differences of these strategies to the fine grid results for the KL probabilities. For comparison, we also include the coarse-mod strategy and some pseudo and centroid strategies that use 20 and 40 clusters and give the ζ difference plots for all these strategies. Transect plot and contour plots for these strategies can also be found at http://depts.washington.edu/ptha/FEMA/.

11 Conclusions and Recommendations

We have implemented the approach summarized in Section 2 and compared several approaches to implementing the steps outlined in the general algorithm presented there.

We conclude with some additional remarks on our findings, the best strategies among those tested, and some additional areas for future research. We have provided additional comments in the sections above; see in particular the notes on clustering strategies found in Section 5.3.

1. The dtopo-clustering of Section H.2 can be done without doing any coarse grid runs. The expense of the centroid strategy that uses this clustering is just the cost of the budget of 20 fine grid runs. A rough approximation can be gotten quite cheaply with this technique, but we recommend it not be depended on as a standalone strategy. Instead, this cheap clustering technique can be used to produce the pseudo-fine runs for the pseudo-dtopo strategy.

2. The etamean-etamax clustering of Section H.1 requires all 400 coarse grid runs be done first. For some probability products, it is better than the dtopo-clustering, and for others it is not. Due to the expense of the 400 coarse grid runs, we do not incur much more cost to produce the remaining 380 pseudo-fine runs for use with a pseudo strategy for better results. Hence, we can not recommend the 20 centroids from this clustering as a standalone strategy.

3. It is possible to improve the coarse grid results to make them more similar to the results that would be obtained by running on a fine grid. This modification procedure is described in Section 7.3.

4. The pseudo strategies from either the dtopo or etamean-etamax clustering require all 400 coarse grid runs be done first. These strategies are considered to give better comparisons to the 400 fine grid results than the centroid only strategies. Sometimes the dtopo clustering works best and sometimes the etamean-etamax clustering works best to make a pseudo strategy. It is interesting that there is very little overlap of the 20 fine grid runs that each clustering strategy yields.

5. When the ζ differences of our strategies (that used a fine grid budget of 20) with the 400 fine grid results were weighted to reflect their conditional probabilities, we found the maximum difference in the ζ values to be around 1.24 (0.91) meters and the mean difference to be very small, around 0.16 (0.12) meters for the pseudo-strategy when the KL (uniform) conditional probabilities were used.

6. A recommendation is to slightly increase the budget of fine grid runs allowed to 10% of the 400, for a total of 40. We did several experiments along these lines. First, we used the KL probabilities and the highest 10 probability runs from each of the four Mw categories (8.6Mw, 8.8Mw, 9.0Mw, and 9.2Mw) for a budget of 40 fine grid runs. The results were extremely good and outperformed the strategies we reported here that used a budget of only 20 fine grid runs as seen in Appendix D. A picture of the zeta-contour for $p = 0.2$ for this strategy can be seen on the cover page of this report. Second, we ran several clustering and pseudo-strategies with this higher budget of 40. Indeed, the results are better than those for the budget of 20, as seen in Appendix E where we give weighted ζ comparisons and plots.

7. For the KL probabilities, Appendix E shows the strategy of using the highest 10 from each of the four Mw categories gave the smallest ζ weighted differences to the fine results with the maximum difference across the rectangular fixed grid being only 0.38 meters and the mean difference being only 0.11 meters. A close second was the pseudo strategy that used the etamean-etamax clustering of the modified coarse grid runs to produce 40 clusters. This strategy gave a maximum difference of 0.65 meters and a mean difference of 0.13 meters. This is interesting, as a strategy using only 40 fine runs was better than any of the pseudo strategies that used 40 fine runs and 360 pseudo runs. This suggests we should consider the probability of each run when we make the clustering. Also, any improvements in the creation of the pseudo-runs would improve the results. These are topics for future work.

8. Section 10 showed that an SVD strategy that used 32 fine grid runs (staying under a budget of 40) gave reasonable results. It is not clear whether increasing the number of fine grid runs to 40 would improve these results. This increase could happen in two ways. First, the number of left singular vectors used could increase from 4 to 5, and still using 2 fine runs per singular vector would yield 10 for each magnitude category, for a total of 40. This might not be successful as the coarse and fine grid runs were not so correlated beyond the first 4 singular vectors. Second, 4 singular vectors could still be used with 3, 3, 2, 2 fine runs associated with them for a total of 10 per magnitude category. More work is needed to see the limitations of this strategy.

9. Another recommended topic for future work is to investigate what budget size is necessary for a cheap clustering strategy that does not require even coarse grid runs (e.g. the dtopo-clustering strategy) to give acceptable results using only the fine grid runs at the centroids. This would have to be compared with the cost of the pseudo-strategies which might give acceptable results with a smaller budget size, but with the additional requirement of the 400 coarse grid runs.

10. This report contains a number of figures, but many more can be viewed online from
    http://depts.washington.edu/ptha/FEMA/
    All of the computer code used in this work has also been archived there, including the code to generate the earthquake realizations, to run the tsunami simulations, to perform the filtering and clustering, and to produce the hazard curves and maps, and to perform comparisons of different approaches.

References

[1] L. Adams, R.J. LeVeque, and F.I. González. The pattern-method for incorporating tidal uncertainty into probabilistic tsunami hazard assessment (ptha). *Natural Hazards*, pages 1–20 (in review), 2014.

[2] Tadashi Annaka, Kenji Satake, Tsutomu Sakakiyama, Ken Yanagisawa, and Nobuo Shuto. Logic-tree Approach for Probabilistic Tsunami Hazard Analysis and its Applications to the Japanese Coasts. *Pure and Applied Geophysics*, 164(2):577–592, March 2007.

[3] M. J. Berger, D. L. George, R. J. LeVeque, and K. T. Mandli. The GeoClaw software for depth-averaged flows with adaptive refinement. *Adv. Water Res.*, 34:1195–1206, 2011.

[4] J. L. Blair, P. A. McCrory, and D. H. Oppenheimer. A geo-referenced 3d model of the Juan de Fuca Slab and associated seismicity. http://pubs.usgs.gov/ds/633/, 2012.

[5] David Burbidge, Phil R Cummins, Richard Mleczko, and Hong Kie Thio. A Probabilistic Tsunami Hazard Assessment for Western Australia. *Pure and Applied Geophysics*, 165(1):2059–2088, December 2008.

[6] Clawpack Development Team. Clawpack software, 2015. Version 5.3.1.

[7] C A CORNELL. Engineering Seismic Risk Analysis. *Bulletin of the Seismological Society of America*, 58(5):1583–1606, 1968.

[8] Hillary Fairbanks, Alireza Doostan, Christian Ketelsen, and Gianluca Iaccarino. A low-rank control variate for multilevel monte carlo simulation of high-dimensional uncertain systems, 2016.

[9] D. Gao. Defining megathrust tsunami sources at northernmost cascadia using thermal and structural information. http://hdl.handle.net/1828/7435, 2016.

[10] Eric L Geist. Local Tsunamis and Earthquake Source Parameters. In Renata Dmowska and Barry Saltzman, editors, *Tsunamigenic Earthquakes and Their Consequences*, pages 117–209. Elsevier, 1998.

[11] Eric L Geist and Tom Parsons. Probabilistic Analysis of Tsunami Hazards*. *Natural Hazards*, 37(3):277–314, 2006.

[12] M. B. Giles. Multilevel Monte Carlo path simulation. *Oper. Res.*, 56:607 – 617, 2008.

[13] F. I. González, E L Geist, B. Jaffe, U Kanoglu, et al. Probabilistic tsunami hazard assessment at Seaside, Oregon, for near-and far-field seismic sources. *J. Geophys. Res.*, 114:C11023, Jan 2009.

[14] F.I. González, R.J. LeVeque, L. Adams, C. Goldfinger, G.R. Priest, and K. Wang. Probabilistic Tsunami Hazard Assessment (PTHA) for Crescent City, CA Final Report September 11,2014. Univ. Washington ResearchWorks Archive: http://hdl.handle.net/1773/25916, 2014.

[15] J L Hammack. A note on tsunamis: their generation and propagation in an ocean of uniform depth. *Journal of Fluid Mechanics*, 60:769–799, 1973.

[16] M Heidarzadeh and A Kijko. A probabilistic tsunami hazard assessment for the Makran subduction zone at the northwestern Indian Ocean. *Natural Hazards*, 2011.

[17] J B Keller. Tsunamis: water waves produced by earthquakes. In Doak C Cox and Ethel McAfee, editors, *IUGG Proceedings of the Tsunami Meeting Associated with the Tenth Pacific Science Congress*, pages 154–166. Proceedings of the Tsunami Meetings Associated with . . . , 1961.

[18] R. J. LeVeque, D. L. George, and M. J. Berger. Tsunami modeling with adaptively refined finite volume methods. *Acta Numerica*, pages 211–289, 2011.

[19] R. J. LeVeque, K. Waagan, F. I. Gonzalez, D. Rim, and G. Lin. Generating random earthquake events for probabilistic tsunami hazard assessment. *Pure Appl. Geophys.*, 2016.

[20] S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, Mar 1982.

[21] D. Melgar, R. J. LeVeque, D. S. Dreger, and R. M. Allen. Kinematic rupture scenarios and synthetic displacement data: An example application to the cascadia subduction zone. *J. Geophys. Res. – Solid Earth*, 121:6658–6674, 2016.

[22] A. Narayan, C. Gittelson, and D. Xiu. A stochastic collocation algorithm with multifidelity models. *SIAM J. Sci. Comput.*, 36(2):A495–A521, 2014.

[23] Yoshimitsu Okada. Surface Deformation Due To Shear And Tensile Faults In A Half-Space. *Bulletin of the Seismological Society of America*, 75:1135–1154, January 1985.

[24] B. Peherstorfer, T. Cui, Y. Marzouk, and K. Willcox. Multifidelity importance sampling. *Comput. Methods Appl. Mech. Engrg.*, 300:490–509, 2016.

[25] M D Petersen, M P Moschetti, P Powers, C S Mueller, K M Haller, A D Frankel, Y Zeng, S Rezaeian, S C Harmsen, O S Boyd, N Field, R Chen, K S Rukstales, N Luco, R L Wheeler, and A Olsen. The 2014 U.S. National Seismic Hazard Maps: A summary of changes to seismic source and ground motion models. In *NCEE 2014 - 10th U.S. National Conference on Earthquake Engineering: Frontiers of Earthquake Engineering*. United States Geological Survey, Reston, United States, January 2014.

[26] Mark D Petersen, Morgan P Moschetti, Peter M Powers, Charles S Mueller, and Kathleen M Haller. Documentation for the 2014 Update of the United States National Seismic Hazard Maps. Technical Report 2014-1091, July 2014.

[27] G. R. Priest, Y. Zhang, R.C. Witter, K. Wang, C. Goldfinger, and L. Stimely. Tsunami impact to Washington and northern Oregon from segment ruptures on the southern Cascadia subduction zone. *Nat. Hazards*, pages 1–22, 2014.

[28] Hong Kie Thio, P G Somerville, and J. Polet. Probabilistic Tsunami Hazard Analysis. In *15th WCEE*, pages 1–9, Lisboa, July 2012.

[29] R. C. Witter, Y. Zhang, K. Wang, G.R. Priest, C. Goldfinger, L. Stimely, J.T. English, and P.A. Ferro. Simulating tsunami inundation for a range of Cascadia megathrust earthquake scenarios at Bandon, Oregon USA. *Geosphere*, 9(6):1783–1803, 2013.

APPENDICES

A Cleaning coarse grid runs

Example results of cleaning coarse grid data to produce the modified coarse grid data are given below. Figure 35 is for the magnitude 9.2 event called run 20. The bottom row shows the coarse data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the coarse and modified coarse data. Notice that $h_{\max}$ is changed in significant places, while eta (η) is not changed at all.

Figure 36 shows how this modified coarse data compares to the corresponding fine data. The bottom row shows the fine data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and modified coarse data. Since the same fine bathymetry is used to produce the modified coarse η and the fine η, the differences of the η values will be the same as the differences of their $h_{\max}$ values as seen in the top row. Together, Figures 35 and 36 show the modified coarse data approximates the fine data much better than the coarse data does.

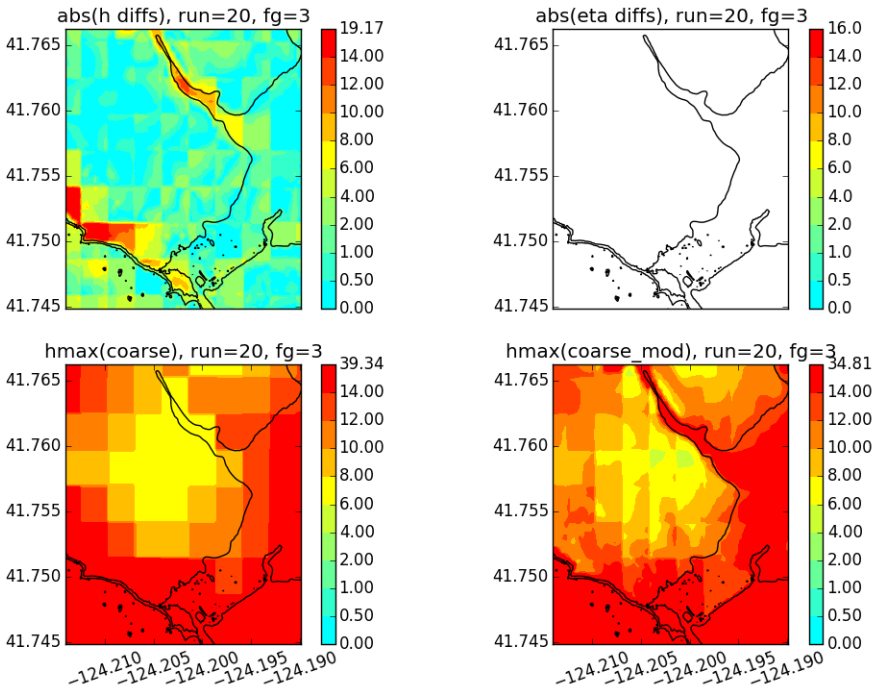

Figure 35: Mw 9.2, run 20: The bottom row shows the coarse data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the coarse and modified coarse data.

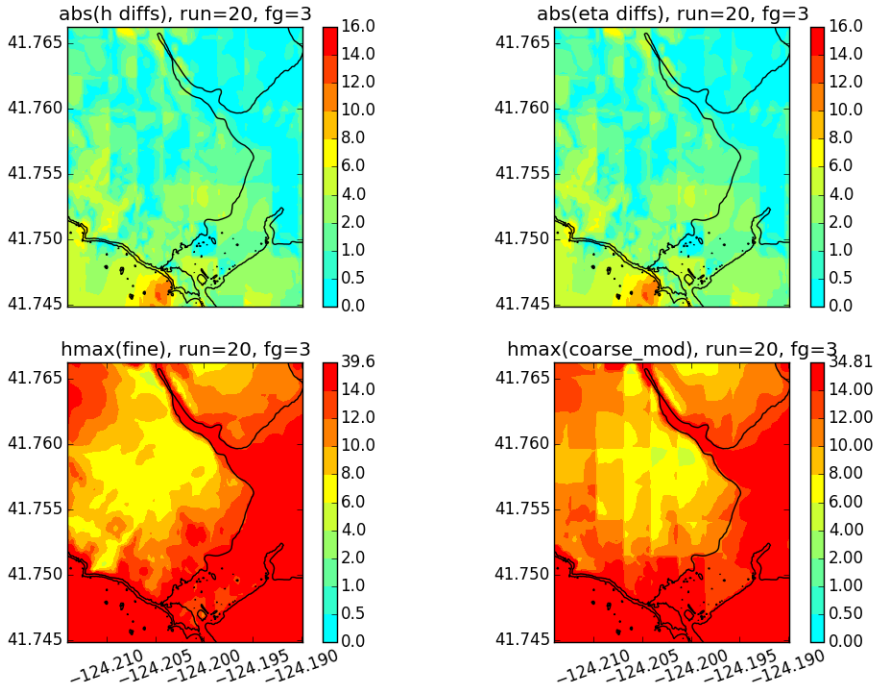

Figure 36: Mw 9.2, run 20: The bottom row shows the fine data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and modified coarse data.

Figure 37 is for the magnitude 8.6 event called run 5. The bottom row shows the coarse data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the coarse and modified coarse data. Notice that for this lower magnitude event, both $h_{\max}$ and η have changed in significant places. It is easy to see that the coarse data had $h_{\max} = 0$ in places because the coarse bathymetry was much higher than the fine bathymetry and no inundation occurred. For these places η was not preserved, nor should it have been. The threshold η value allowed inundation at these places in the modified coarse data.

Figure 38 shows how this modified coarse data compares to the corresponding fine data. The bottom row shows the fine data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and modified coarse data. Since the same fine bathymetry is used to produce the modified coarse η and the fine η, the differences of the η values will be the same as the differences of their $h_{\max}$ values as seen in the top row. Together, Figures 37 and 38 show the modified coarse data approximates the fine data much better than the coarse data does.

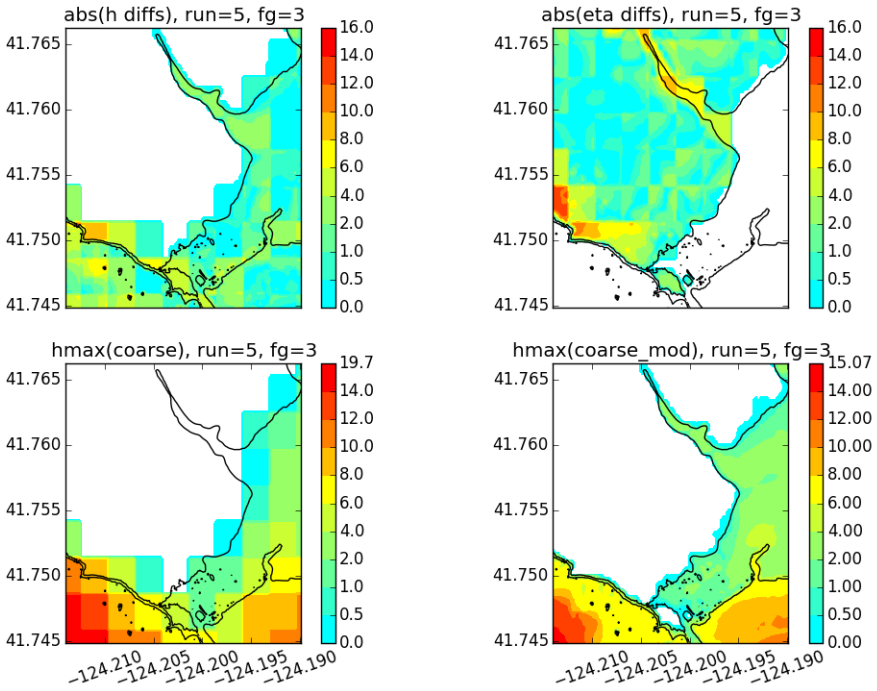

Figure 37: Mw 8.6, run 5: The bottom row shows the coarse data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the coarse and modified coarse data.

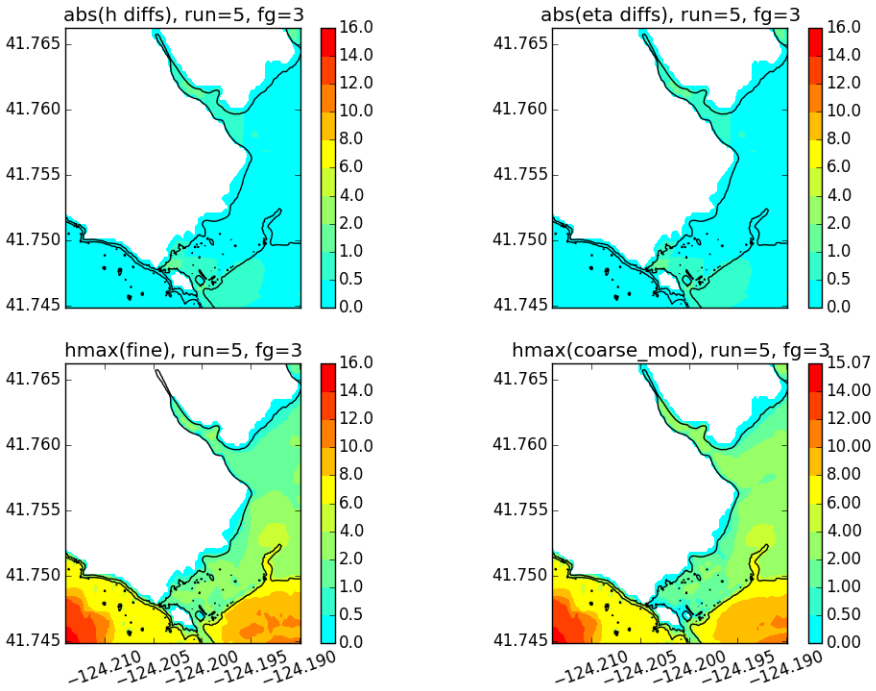

Figure 38: Mw 8.6, run 5: The bottom row shows the fine data and the modified coarse data. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and modified coarse data.

B Pseudo fine grid runs

Example results comparing the pseudo fine grid data to the fine grid data is given in Figure 39 for the magnitude 8.6 event called run 5 when the etamean-etamax clustering strategy was used. The bottom row shows the fine data and its pseudo fine approximation. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and pseudo fine data. These differences are the same since the fine bathymetry is used to define η in both cases.

Figure 40 shows the same information, but when the dtopo clustering strategy is used. Together, Figures 39 and 40 show both clustering strategies yield pseudo fine data that approximates the fine data fairly well.

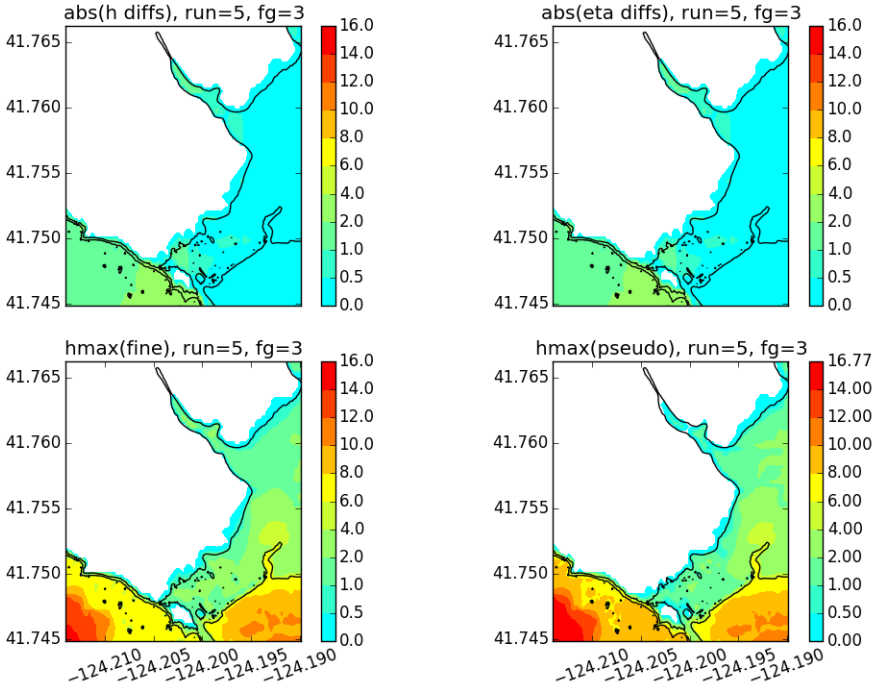

Figure 39: Mw 8.6, run 5: The bottom row shows the fine data and its pseudo fine approximation when etamean-etamax clustering is used. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and pseudo fine data.

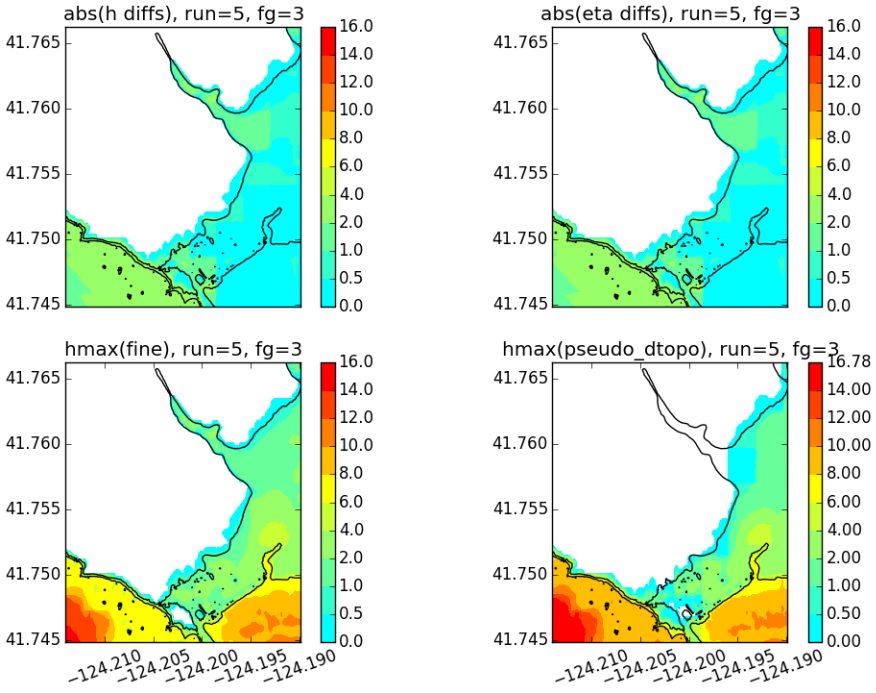

Figure 40: Mw 8.6, run 5: The bottom row shows the fine data and its pseudo fine approximation when dtopo clustering is used. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and pseudo fine data.

Figure 41 is for the magnitude 9.2 event called run 20 when the etamean-etamax clustering is used. The bottom row shows the fine data and its pseudo fine approximation. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and pseudo fine data. For this event, the pseudo fine approximation is perfect because this event was one of the centroids of the etamean-etamax clustering.

Figure 42 shows the same information, but when the dtopo clustering strategy is used. This event was not a centroid of the dtopo clustering. Together, Figures 41 and 42 show both clustering strategies yield pseudo fine data that approximates the fine data fairly well.

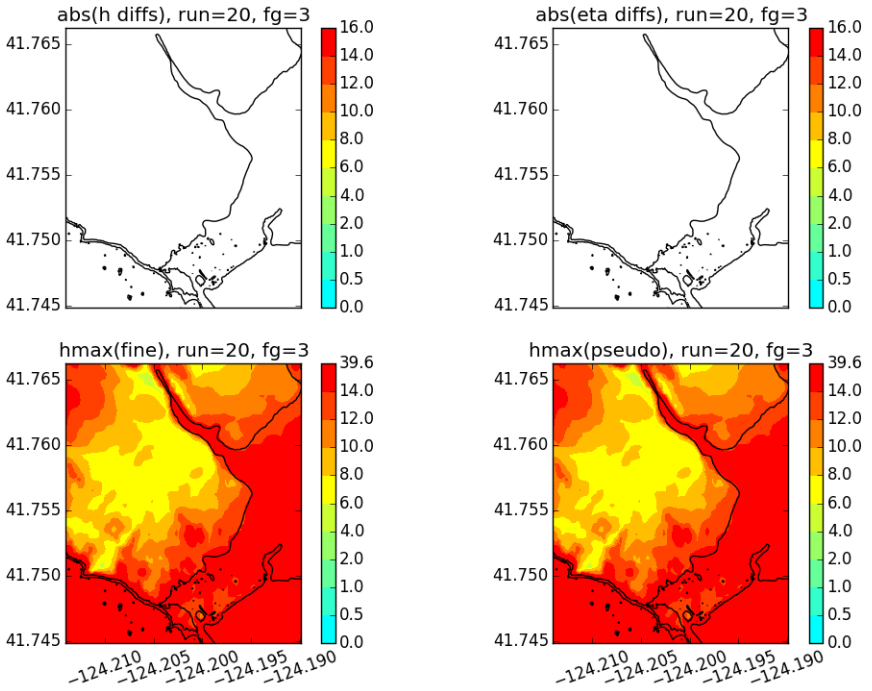

Figure 41: Mw 9.2, run 20: The bottom row shows the fine data and its pseudo fine approximation when etamean-etamax clustering is used. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and pseudo fine data.

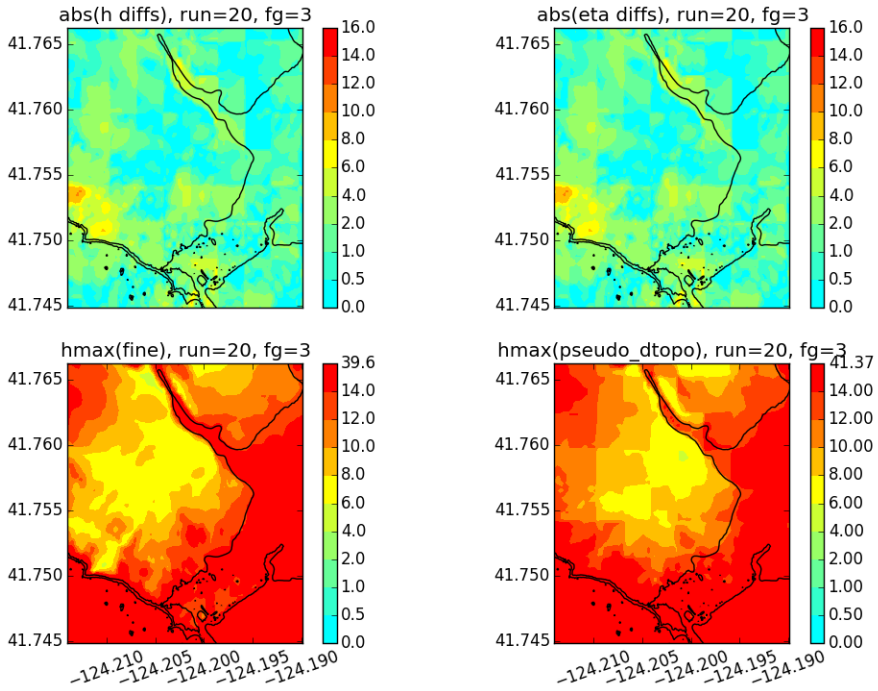

Figure 42: Mw 9.2, run 20: The bottom row shows the fine data and its pseudo fine approximation when dtopo clustering is used. The top row gives the differences of the $h_{\max}$ values and the eta (η) values of the fine and pseudo fine data.

C SVD strategy in detail

Here we outline the complete details of the SVD strategy shown in Section 7.6. This strategy is carried out separately for each of the four magnitudes represented in the study, recalling there are 100 fine-grid runs and 100 coarse-grid runs for each magnitude.

For each of the rectangular fgmax data, let us flatten the 88×78 grid data of $h_{\max} + B0$ values into column vectors, then stack them into a 6864×100 matrix. Let us denote by A_f such a matrix formed for the 100 fine-grid runs and by A_c a corresponding matrix for the 100 coarse-grid runs. We compute the SVD of these two matrices,

$$\begin{aligned} A_c &= U_c \Sigma_c V_c^T = U_c Y_c \quad \text{where } Y_c = \Sigma_c V_c^T, \\ A_f &= U_f \Sigma_f V_f^T = U_f Y_f \quad \text{where } Y_f = \Sigma_f V_f^T. \end{aligned} \tag{4}$$

This strategy relies on the fact that first columns of U_c and those of U_f are similar, as well as the first rows of Y_c and those of Y_f. The similarity between U_c and U_f, for example, are shown in Figure 43.

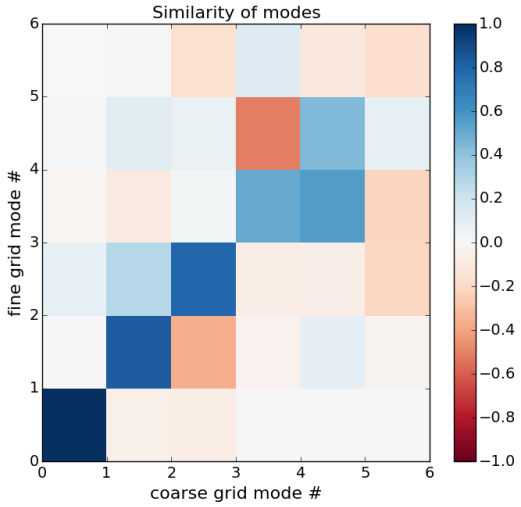

Figure 43: High similarity between coarse and fine singular modes, for the first 6 modes taken from magnitude 9.0 events. Similarity is measured by the cosine of the angle between the two unit vectors.

The values on the diagonal of $U_c^T U_f$ have absolute values close to 1, indicating a strong similarity. The similarity begins to break down after the 4th to 5th mode, so we use the first 4 modes to construct an approximation to A_f.

In our study, computing A_f would require conducting all the fine-grid runs. Therefore, in order to compute an approximation to A_f, which we denote by $\hat{A}_f$, we attempt to construct an approximation to the first four columns of U_f and the first four rows of Y_f using minimal amount of information from the fine-grid runs. Since the coarse and fine runs are similar for the first four modes, we can conclude that the coarse runs that have large components (values in the first 4 rows of Y_c) in the first four left singular directions can be used to build an approximation to the first 4 columns of U_f which we call $\hat{U}_f$. Since we are given Y_c, we sort the columns (indexed by i) of its first four rows 4 ways for $j = 1, 2, 3, 4$, largest to smallest by the value

$$\frac{|Y_{c,ji}|}{\sqrt{\sum_{\substack{k \neq j \\ k \in \{1,2,3,4\}}} |Y_{c,ki}|^2}}. \tag{5}$$

In this study, the largest 2 columns for each of $j = 1, 2, 3, 4$ were chosen, so that we arrive at 8 column indices of the coarse grid runs. These indices will be the fine-grid runs we will perform, and we will denote the set of these indices by I.

Taking the fine-grid rectangular fgmax $h_{\max} + B0$ values from the runs indexed by I, we again flatten the grid data and form a 6864×8 matrix, which we denote by B_f. The SVD of B_f is

$$B_f = \bar{U}_f \bar{\Sigma}_f \bar{V}_f^T = \bar{U}_f \bar{Y}_f \quad \text{where } \bar{Y}_f = \bar{\Sigma}_f \bar{V}_f^T. \tag{6}$$

We can now use the first four columns of $\bar{U}_f$ as our value for $\hat{U}_f$.

Now, we turn to the approximation of the first four rows of Y_f. While the first four rows of Y_c and Y_f are similar, we can not simply equate them. We could use the known information in the first four rows of $\bar{Y}_f$ (called $\tilde{Y}_f$) to find a linear map T that can transform the first four rows of Y_c (called $\bar{Y}_c$) into its fine-grid counterpart ($\tilde{Y}_f$) in the least-squares sense, by solving the minimization problem

$$\text{minimize} \|T\bar{Y}_{c,i} - \tilde{Y}_{f,i}\|^2 \quad \text{for } i \in I. \tag{7}$$

The components of the computed T shown in Figure 44 reveal that it is roughly a diagonal map, indicating that there is also a strong similarity between the columns of $\bar{Y}_c$ and $\tilde{Y}_f$ represented in I.

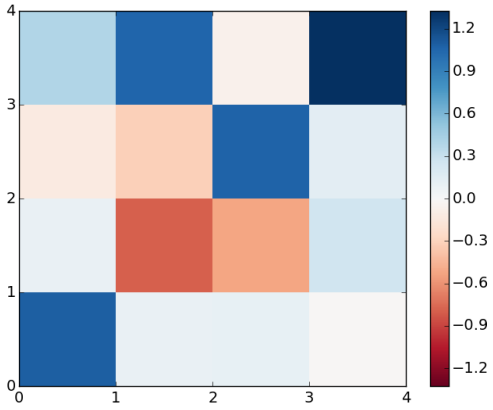

Figure 44: Components of T in the minimization problem of (7).

Finally, we choose $\hat{A}_f = \hat{U}_f T \bar{Y}_c$ to be the approximation to A_f.

D Strategy differences for KL probabilities

```
Mean and max differences to all 400 fine for the following:
 Truth:    scenario_all_fine                (finding differences to this one)
 Strategy: scenario_all_pseudo              (etamean-etamax clustering)
 Strategy: scenario_dzCC_dtopo_etamax_20cls (dtopo clustering)
 Strategy: scenario_coarse_eta_20cls        (etamean-etamax clustering)
 Strategy: scenario_all_pseudo_dtopo        (dtopo clustering)
 Strategy: scenario_all_SVD_40              (SVD, 32 fine, 8 each magnitude)
 Strategy: scenario_highest10_each          (** budget of 40 fine runs)

p-contour differences for zeta= 0.00
 mean and max for scenario_all_pseudo was                0.03325 0.27991
 mean and max for scenario_dzCC_dtopo_etamax_20cls was   0.04130 0.35018
 mean and max for scenario_coarse_eta_20cls was          0.03921 0.21411
 mean and max for scenario_all_pseudo_dtopo was          0.04903 0.34712
 mean and max for scenario_all_SVD_40 was                0.06869 0.41052
 mean and max for scenario_highest10_each was            0.01353 0.05596

p-contour differences for zeta= 1.00
 mean and max for scenario_all_pseudo was                0.02581 0.19460
 mean and max for scenario_dzCC_dtopo_etamax_20cls was   0.04826 0.34889
 mean and max for scenario_coarse_eta_20cls was          0.03355 0.18312
 mean and max for scenario_all_pseudo_dtopo was          0.03998 0.34330
 mean and max for scenario_all_SVD_40 was                0.03921 0.20712
 mean and max for scenario_highest10_each was            0.01635 0.05676

p-contour differences for zeta= 2.00
 mean and max for scenario_all_pseudo was                0.01837 0.15492
 mean and max for scenario_dzCC_dtopo_etamax_20cls was   0.05147 0.35724
 mean and max for scenario_coarse_eta_20cls was          0.03209 0.18889
 mean and max for scenario_all_pseudo_dtopo was          0.02564 0.33681
 mean and max for scenario_all_SVD_40 was                0.01883 0.39885
 mean and max for scenario_highest10_each was            0.01750 0.05785

p-contour differences for zeta= 3.00
 mean and max for scenario_all_pseudo was                0.02053 0.13470
 mean and max for scenario_dzCC_dtopo_etamax_20cls was   0.05592 0.33386
 mean and max for scenario_coarse_eta_20cls was          0.03504 0.17133
 mean and max for scenario_all_pseudo_dtopo was          0.02525 0.27915
 mean and max for scenario_all_SVD_40 was                0.02028 0.39522
 mean and max for scenario_highest10_each was            0.01784 0.05800

p-contour differences for zeta= 4.00
 mean and max for scenario_all_pseudo was                0.01920 0.15859
 mean and max for scenario_dzCC_dtopo_etamax_20cls was   0.05796 0.32510
 mean and max for scenario_coarse_eta_20cls was          0.03486 0.15254
 mean and max for scenario_all_pseudo_dtopo was          0.02844 0.22041
 mean and max for scenario_all_SVD_40 was                0.03302 0.34902
 mean and max for scenario_highest10_each was            0.01871 0.06474
```

```
p-contour differences for zeta= 6.00
 mean and max for scenario_all_pseudo was                 0.01581 0.12986
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.06601 0.42750
 mean and max for scenario_coarse_eta_20cls was           0.03096 0.21138
 mean and max for scenario_all_pseudo_dtopo was           0.02145 0.15163
 mean and max for scenario_all_SVD_40 was                 0.02715 0.23510
 mean and max for scenario_highest10_each was             0.01328 0.05597

p-contour differences for zeta= 8.00
 mean and max for scenario_all_pseudo was                 0.01781 0.14164
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.04400 0.31759
 mean and max for scenario_coarse_eta_20cls was           0.02520 0.14631
 mean and max for scenario_all_pseudo_dtopo was           0.01326 0.10691
 mean and max for scenario_all_SVD_40 was                 0.02369 0.19362
 mean and max for scenario_highest10_each was             0.01379 0.04740

p-contour differences for zeta= 10.00
 mean and max for scenario_all_pseudo was                 0.01053 0.14561
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.03086 0.16069
 mean and max for scenario_coarse_eta_20cls was           0.01451 0.14145
 mean and max for scenario_all_pseudo_dtopo was           0.01031 0.09573
 mean and max for scenario_all_SVD_40 was                 0.01735 0.17246
 mean and max for scenario_highest10_each was             0.01361 0.05465

p-contour differences for zeta= 12.00
 mean and max for scenario_all_pseudo was                 0.00742 0.14041
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.02969 0.20001
 mean and max for scenario_coarse_eta_20cls was           0.01430 0.14243
 mean and max for scenario_all_pseudo_dtopo was           0.00736 0.10113
 mean and max for scenario_all_SVD_40 was                 0.01483 0.16029
 mean and max for scenario_highest10_each was             0.01053 0.04759

zeta-contour differences for p=0.1
 mean and max for scenario_all_pseudo was                 0.65543 4.79177
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.86499 3.85224
 mean and max for scenario_coarse_eta_20cls was           0.73280 5.52382
 mean and max for scenario_all_pseudo_dtopo was           0.36183 1.89197
 mean and max for scenario_all_SVD_40 was                 0.80124 2.93863
 mean and max for scenario_highest10_each was             1.04443 2.89255

zeta-contour differences for p=0.2
 mean and max for scenario_all_pseudo was                 0.50798 3.47264
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    1.08555 3.40863
 mean and max for scenario_coarse_eta_20cls was           1.09745 3.81849
 mean and max for scenario_all_pseudo_dtopo was           0.67635 3.44487
 mean and max for scenario_all_SVD_40 was                 0.57404 2.89988
 mean and max for scenario_highest10_each was             0.24138 1.67055
```

```
zeta-contour differences for p=0.3
 mean and max for scenario_all_pseudo was                0.35866 1.85071
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.61153 2.88502
 mean and max for scenario_coarse_eta_20cls was          0.61011 2.62107
 mean and max for scenario_all_pseudo_dtopo was          0.46467 2.87974
 mean and max for scenario_all_SVD_40 was                0.44390 2.77637
 mean and max for scenario_highest10_each was            0.81164 2.70336

zeta-contour differences for p=0.4
 mean and max for scenario_all_pseudo was                0.71101 3.21702
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.97569 4.06048
 mean and max for scenario_coarse_eta_20cls was          1.35825 3.82501
 mean and max for scenario_all_pseudo_dtopo was          0.45732 2.84070
 mean and max for scenario_all_SVD_40 was                0.52500 2.68276
 mean and max for scenario_highest10_each was            0.05384 0.84205

zeta-contour differences for p=0.5
 mean and max for scenario_all_pseudo was                0.46703 2.67963
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.20411 1.39488
 mean and max for scenario_coarse_eta_20cls was          0.54395 2.75662
 mean and max for scenario_all_pseudo_dtopo was          0.26784 1.91235
 mean and max for scenario_all_SVD_40 was                0.38765 2.48476
 mean and max for scenario_highest10_each was            0.19839 1.04073

zeta-contour differences for p=0.6
 mean and max for scenario_all_pseudo was                0.36037 1.95435
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    0.79710 1.83306
 mean and max for scenario_coarse_eta_20cls was          0.57318 2.18430
 mean and max for scenario_all_pseudo_dtopo was          0.28055 1.75859
 mean and max for scenario_all_SVD_40 was                0.56798 2.83285
 mean and max for scenario_highest10_each was            0.20575 1.66151

zeta-contour differences for p=0.7
 mean and max for scenario_all_pseudo was                0.21960 0.93398
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    1.10822 2.83315
 mean and max for scenario_coarse_eta_20cls was          0.27746 1.27390
 mean and max for scenario_all_pseudo_dtopo was          0.45553 2.74874
 mean and max for scenario_all_SVD_40 was                0.61447 2.76804
 mean and max for scenario_highest10_each was            0.06244 0.52286

zeta-contour differences for p=0.8
 mean and max for scenario_all_pseudo was                0.23294 1.55166
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    1.22250 3.43998
 mean and max for scenario_coarse_eta_20cls was          0.31104 1.65237
 mean and max for scenario_all_pseudo_dtopo was          0.59925 3.51161
 mean and max for scenario_all_SVD_40 was                0.40445 2.70863
 mean and max for scenario_highest10_each was            0.12507 0.62055
```

```
zeta-contour differences for p=0.9
 mean and max for scenario_all_pseudo was                0.15953 1.53221
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    1.14875 3.05149
 mean and max for scenario_coarse_eta_20cls was           0.63211 2.13293
 mean and max for scenario_all_pseudo_dtopo was           0.70728 3.06078
 mean and max for scenario_all_SVD_40 was                 0.29474 2.28850
 mean and max for scenario_highest10_each was             0.25633 0.93414

zeta-contour differences for p=0.95
 mean and max for scenario_all_pseudo was                0.37055 4.47455
 mean and max for scenario_dzCC_dtopo_etamax_20cls was    1.62199 6.70065
 mean and max for scenario_coarse_eta_20cls was           0.98401 4.47455
 mean and max for scenario_all_pseudo_dtopo was           1.46288 8.05844
 mean and max for scenario_all_SVD_40 was                 0.25358 2.28354
 mean and max for scenario_highest10_each was             0.07442 0.34842
```

E Strategy weighted zeta differences for KL probabilities

```
Mean and max zeta-contour differences to all 400 fine for the following:

 Truth:    scenario_all_fine                (finding differences to this one)
 Strategy: scenario_all_coarse_mod          (the 400 coarse_mod)

       20-Centroid Strategies
 Strategy: scenario_dzCC_PE_20cls           (dtopo-pe clustering of coarse_mod runs)
 Strategy: scenario_coarse_eta_20cls        (etamean-etamax clustering of coarse_runs)
 Strategy: scenario_coarse_etaCM_20cls      (etamean-etamax clustering of coarse_mod runs)
 Strategy: scenario_dzCC_dtopo_etamax_20cls (dtopo clustering of coarse_runs)

       Pseudo-Strategies using 20 fine
 Strategy: scenario_all_pseudo              (etamean-etamax clustering of coarse_runs)
 Strategy: scenario_all_pseudoCM            (etamean-etamax clustering of coarse_mod runs)
 Strategy: scenario_all_pseudo_dtopo        (dtopo clustering of coarse_runs)
 Strategy: scenario_all_pseudoCM_dzCC_PE    (dtopo-pe clustering of coarse_mod runs)

       40-Centroid Strategies
 Strategy: scenario_highest10_each          (** budget of 40 fine runs)
 Strategy: scenario_coarse_eta_40cls        (etamean-etamax clustering of coarse_runs)
 Strategy: scenario_coarse_etaCM_40cls      (etamean-etamax clustering of coarse_mod runs)
 Strategy: scenario_dtopo_40cls             (dtopo clustering of coarse_runs)
 Strategy: scenario_all_pseudoCM_dzCC_PE_40 (dtopo-pe clustering of coarse_mod runs)

       Pseudo-Strategies using 40 fine
 Strategy: scenario_all_pseudo_40           (etamean-etamax clustering of coarse_runs)
 Strategy: scenario_all_pseudoCM_40         (etamean-etamax clustering of coarse_mod runs)
 Strategy: scenario_all_pseudo_dtopo_40     (dtopo clustering of coarse_runs)
 Strategy: scenario_all_pseudoCM_dzCC_PE_40 (dtopo-pe clustering of coarse_mod runs)
 Strategy: scenario_all_SVD_40              (SVDs on coarse, used 32 fine runs, 8 each magnitude)

 mean and max for scenario_all_coarse_mod was            0.38581 2.46464
 mean and max for scenario_dzCC_PE_20cls was             0.30763 1.21847
 mean and max for scenario_coarse_eta_20cls was          0.26644 1.34041
 mean and max for scenario_coarse_etaCM_20cls was        0.26651 1.20747
 mean and max for scenario_dzCC_dtopo_etamax_20cls was   0.45621 2.43414
 mean and max for scenario_all_pseudo was                0.16243 1.23804
 mean and max for scenario_all_pseudoCM was              0.16729 1.16061
 mean and max for scenario_all_pseudo_dtopo was          0.24442 1.70556
 mean and max for scenario_all_pseudoCM_dzCC_PE was      0.18586 0.82847
 mean and max for scenario_highest10_each was            0.10635 0.37572
 mean and max for scenario_dzCC_PE_40cls was             0.36488 1.94609
 mean and max for scenario_coarse_eta_40cls was          0.19975 1.21092
 mean and max for scenario_coarse_etaCM_40cls was        0.21045 0.98994
 mean and max for scenario_dtopo_40cls was               0.29449 1.19787
 mean and max for scenario_all_pseudo_40 was             0.13494 0.88091
 mean and max for scenario_all_pseudoCM_40 was           0.13557 0.65425
 mean and max for scenario_all_pseudo_dtopo_40 was       0.14392 1.01533
 mean and max for scenario_all_pseudoCM_dzCC_PE_40 was   0.22660 1.59411
 mean and max for scenario_all_SVD_40 was                0.22602 2.48940
```

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 45 were (2.46,0.39) for scenario-all-coarse-mod, (1.22,0.31) for scenario-dzCC-PE-20cls, (2.43,0.46) for scenario-dzCC-dtopo-etamax-20cls, (1.34,0.27) for scenario-coarse-eta-20cls, and (1.21,0.27) for scenario-coarse-etaCM-20cls.

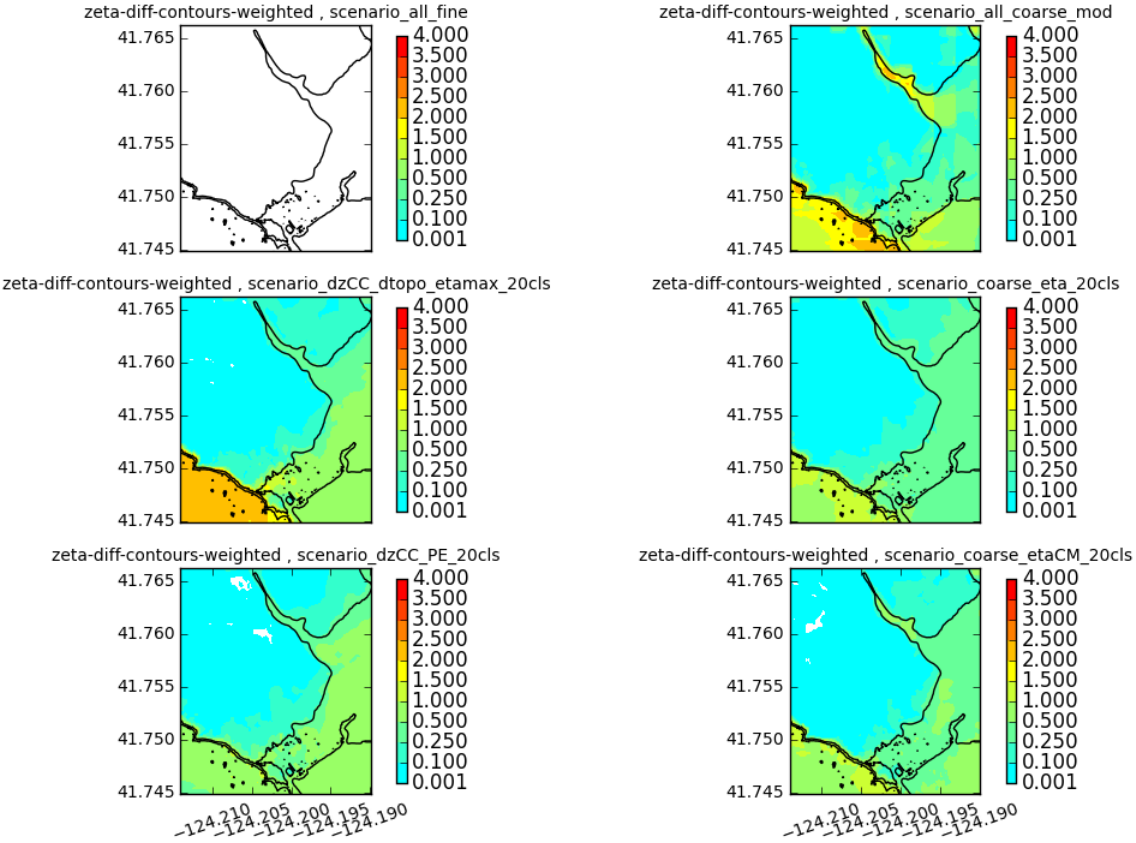

Figure 45: Weighted zeta difference contours of coarse-mod and centroid strategies using 20 clusters and KL probabilities

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 46 were (2.46,0.39) for scenario-all-coarse-mod, (1.24,0.16) for scenario-all-pseudo, (1.16,0.17) for scenario-all-pseudoCM, (1.71,0.24) for scenario-all-pseudo-dtopo, and (0.83,0.19) for scenario-all-pseudoCM-dzCC-PE.

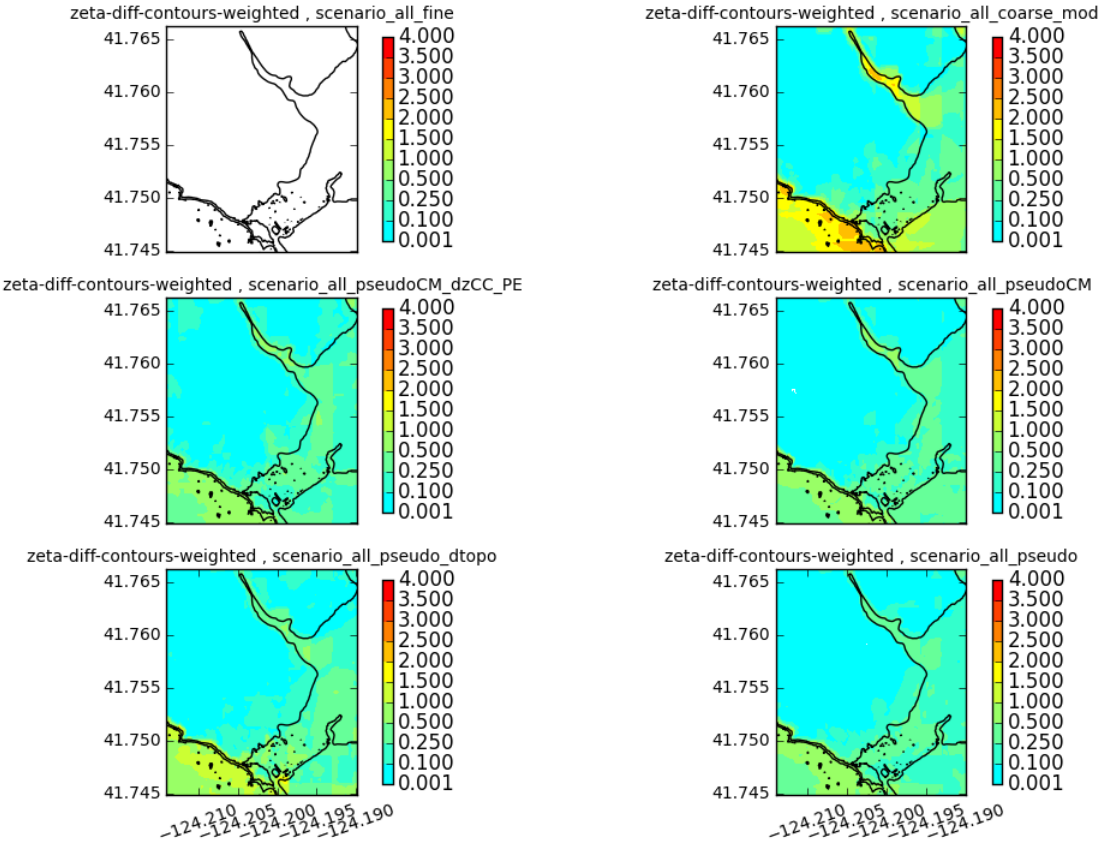

Figure 46: Weighted zeta difference contours of coarse-mod and pseudo strategies using 20 clusters and KL probabilities

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 47 were (0.38,0.11) for scenario-highest10-each, (1.21,0.20) for scenario-coarse-eta-40cls, (1.20,0.29) for scenario-dtopo-40cls, (0.88,0.13) for scenario-all-pseudo-40, and (1.02,0.14) for scenario-all-pseudo-dtopo-40.

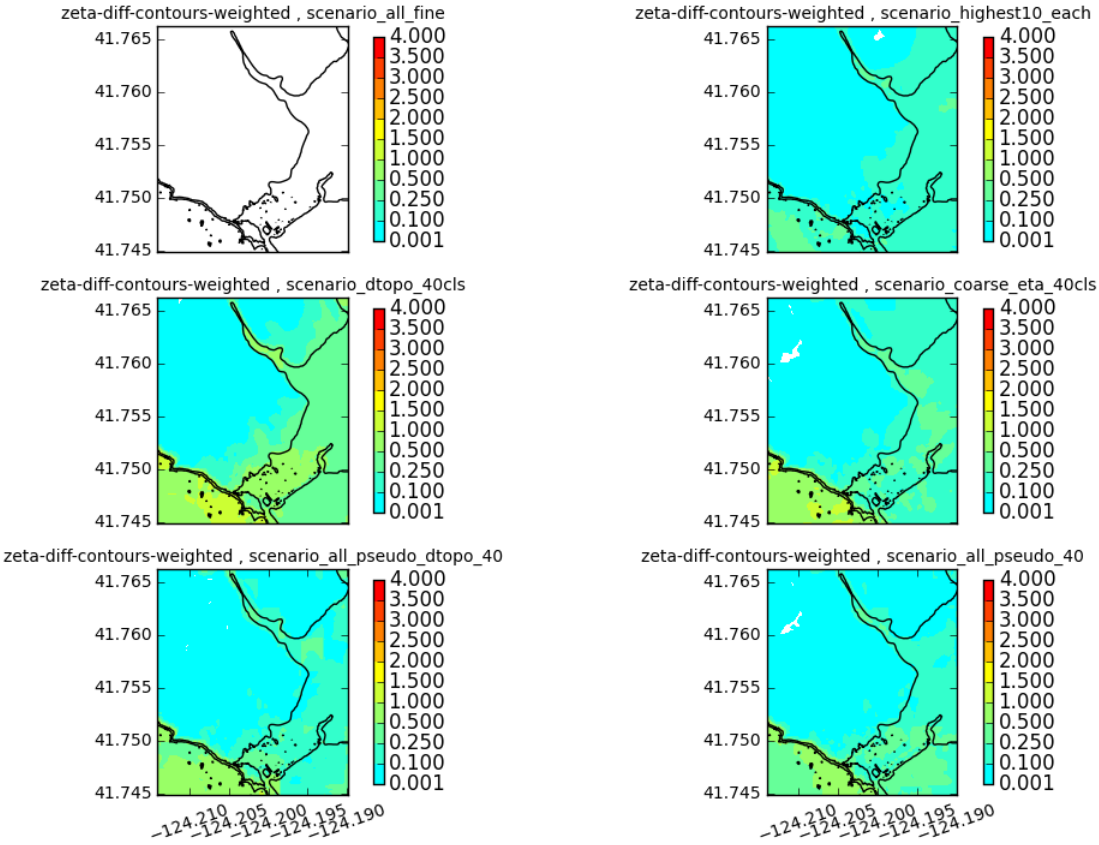

Figure 47: Weighted zeta difference contours of centroid and pseudo strategies using 40 clusters and KL probabilities

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 48 were (0.38,0.11) for scenario-highest10-each, (0.99,0.21) for scenario-coarse-etaCM-40cls, (1.95,0.36) for scenario-dzCC-PE-40cls, (0.65,0.13) for scenario-all-pseudoCM-40, and (1.59,0.23) for scenario-all-pseudoCM-dzCC-PE-40.

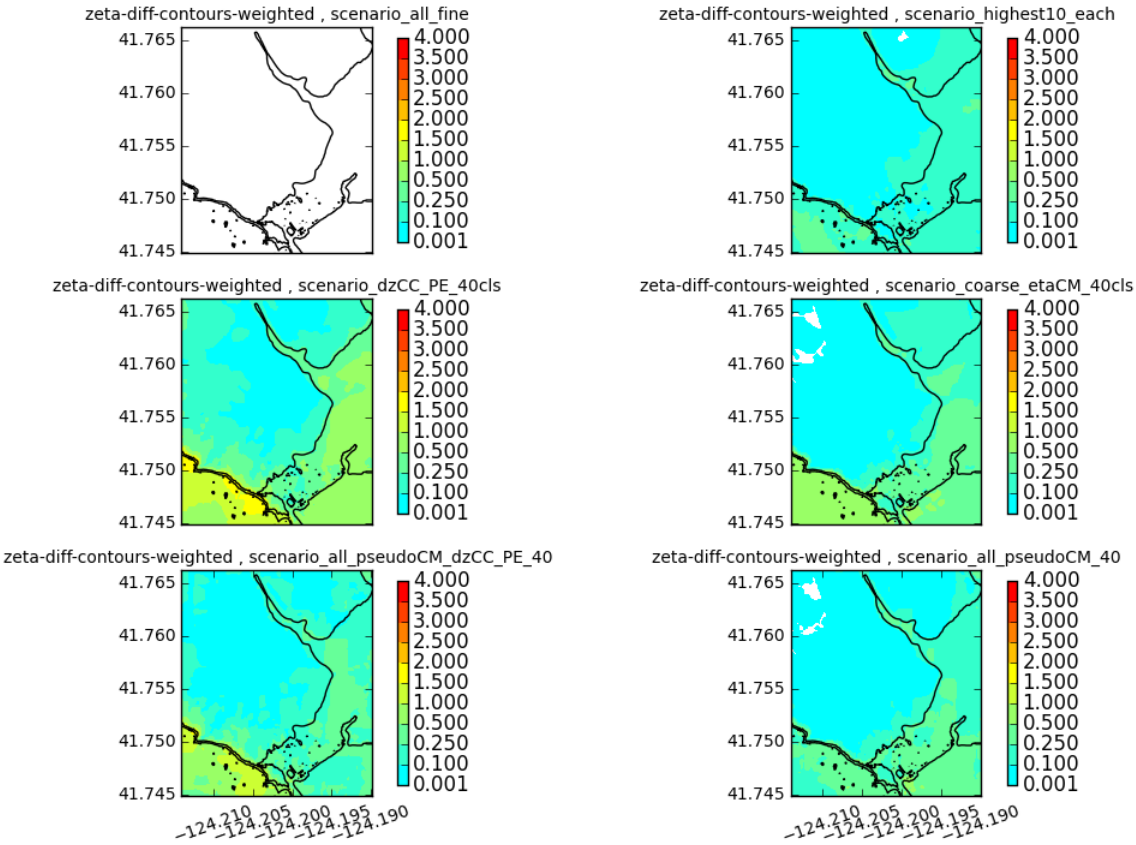

Figure 48: Weighted zeta difference contours of centroid and pseudo strategies using 40 clusters and KL probabilities

The (maximum, mean) of the weighted value of the ζ differences across the fixed grid in Figure 49 were (0.38,0.11) for scenario-highest10-each, (0.65,0.13) for scenario-all-pseudoCM-40, and (2.49,0.23) for scenario-all-SVD-40.

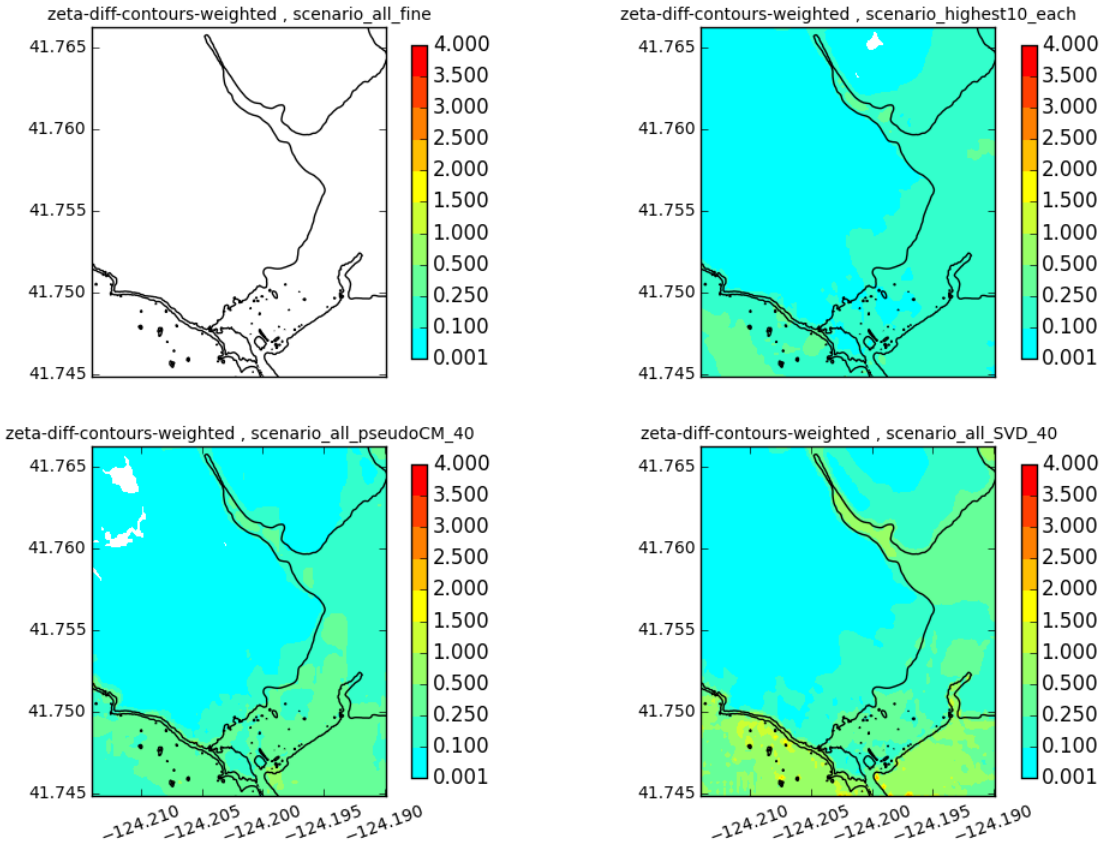

Figure 49: Weighted zeta difference contours for scenario-highest10-each, scenario-all-pseudoCM-40 and scenario-all-SVD-40 using KL probabilities

F Strategy differences for uniform probabilities

```
Mean and max differences to all 400 fine for the following:
 Truth:    scenario_all_fine_uniform         (finding differences to this one)
 Strategy: scenario_all_pseudo_uniform       (etamean-etamax clustering)
 Strategy: scenario_dzCC_dtopo_20cls_uniform (dtopo clustering)
 Strategy: scenario_c_eta_20cls_uniform      (etamean-etamax clustering)
 Strategy: scenario_all_pseudo_dtopo_uniform (dtopo clustering)

p-contour differences for zeta= 0.00
 mean and max for scenario_all_pseudo_uniform was         0.02645 0.21000
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.04627 0.36500
 mean and max for scenario_c_eta_20cls_uniform was        0.02992 0.22200
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.05302 0.35300

p-contour differences for zeta= 1.00
 mean and max for scenario_all_pseudo_uniform was         0.02012 0.14500
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.04608 0.35300
 mean and max for scenario_c_eta_20cls_uniform was        0.03043 0.18200
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.03523 0.34800

p-contour differences for zeta= 2.00
 mean and max for scenario_all_pseudo_uniform was         0.01702 0.14200
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.05268 0.36700
 mean and max for scenario_c_eta_20cls_uniform was        0.03337 0.18500
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.02379 0.33700

p-contour differences for zeta= 3.00
 mean and max for scenario_all_pseudo_uniform was         0.01731 0.14900
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.05761 0.33100
 mean and max for scenario_c_eta_20cls_uniform was        0.03464 0.18100
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.02392 0.32000

p-contour differences for zeta= 4.00
 mean and max for scenario_all_pseudo_uniform was         0.01478 0.17000
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.05846 0.31900
 mean and max for scenario_c_eta_20cls_uniform was        0.03315 0.16700
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.02553 0.21500

p-contour differences for zeta= 6.00
 mean and max for scenario_all_pseudo_uniform was         0.01273 0.12500
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.06300 0.41200
 mean and max for scenario_c_eta_20cls_uniform was        0.02656 0.16000
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.01964 0.15800

p-contour differences for zeta= 8.00
 mean and max for scenario_all_pseudo_uniform was         0.01682 0.12500
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.03869 0.34500
 mean and max for scenario_c_eta_20cls_uniform was        0.02332 0.15200
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.01079 0.06900
```

```
p-contour differences for zeta= 10.00
 mean and max for scenario_all_pseudo_uniform was         0.00718 0.08200
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.02840 0.16400
 mean and max for scenario_c_eta_20cls_uniform was        0.01332 0.09800
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.00903 0.07000

p-contour differences for zeta= 12.00
 mean and max for scenario_all_pseudo_uniform was         0.00716 0.08000
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.02768 0.17300
 mean and max for scenario_c_eta_20cls_uniform was        0.01347 0.09800
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.00568 0.04800

zeta-contour differences for p=0.1
 mean and max for scenario_all_pseudo_uniform was         0.74908 5.40000
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   1.42215 3.35880
 mean and max for scenario_c_eta_20cls_uniform was        1.02259 5.99672
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.37461 2.53000

zeta-contour differences for p=0.2
 mean and max for scenario_all_pseudo_uniform was         0.33113 1.75667
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.94272 2.41488
 mean and max for scenario_c_eta_20cls_uniform was        0.73354 2.17523
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.38990 2.14000

zeta-contour differences for p=0.3
 mean and max for scenario_all_pseudo_uniform was         0.34151 2.12667
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   1.59296 3.31657
 mean and max for scenario_c_eta_20cls_uniform was        0.59284 2.81619
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.32019 1.87000

zeta-contour differences for p=0.4
 mean and max for scenario_all_pseudo_uniform was         0.24029 2.00667
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   1.07266 4.56467
 mean and max for scenario_c_eta_20cls_uniform was        0.64133 2.18188
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.49245 3.87045

zeta-contour differences for p=0.5
 mean and max for scenario_all_pseudo_uniform was         0.33919 2.06362
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.41259 1.58086
 mean and max for scenario_c_eta_20cls_uniform was        0.29292 1.61868
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.27056 1.76000

zeta-contour differences for p=0.6
 mean and max for scenario_all_pseudo_uniform was         0.24477 1.33333
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   0.88467 1.65984
 mean and max for scenario_c_eta_20cls_uniform was        0.46042 1.55212
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.32231 1.89000
```

```
zeta-contour differences for p=0.7
 mean and max for scenario_all_pseudo_uniform was         0.18283 0.98333
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   1.13806 3.26836
 mean and max for scenario_c_eta_20cls_uniform was        0.48335 1.38626
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.51376 3.25149

zeta-contour differences for p=0.8
 mean and max for scenario_all_pseudo_uniform was         0.19893 1.04000
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   1.25523 3.92343
 mean and max for scenario_c_eta_20cls_uniform was        0.33307 0.94725
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.68502 3.51135

zeta-contour differences for p=0.9
 mean and max for scenario_all_pseudo_uniform was         0.20132 1.25167
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   1.41365 3.63170
 mean and max for scenario_c_eta_20cls_uniform was        0.42855 1.32959
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.80430 2.99532

zeta-contour differences for p=0.95
 mean and max for scenario_all_pseudo_uniform was         0.38230 1.51333
 mean and max for scenario_dzCC_dtopo_20cls_uniform was   1.80327 4.24163
 mean and max for scenario_c_eta_20cls_uniform was        0.42427 0.95634
 mean and max for scenario_all_pseudo_dtopo_uniform was   0.83514 2.92024
```

G Generating earthquake realizations

We use a Karhunen-Loève (K-L) expansion to generate a random slip pattern on a specified fault geometry, sampled from random fields with a specified covariance structure. The general approach we use is described in more detail in [19].

For the fault geometry, we used a subset of the Cascadia Subduction Zone, taking only the southern-most 8 fault segments as illustrated in Figure 50. The same geometry was used in [19] and is discussed further there.

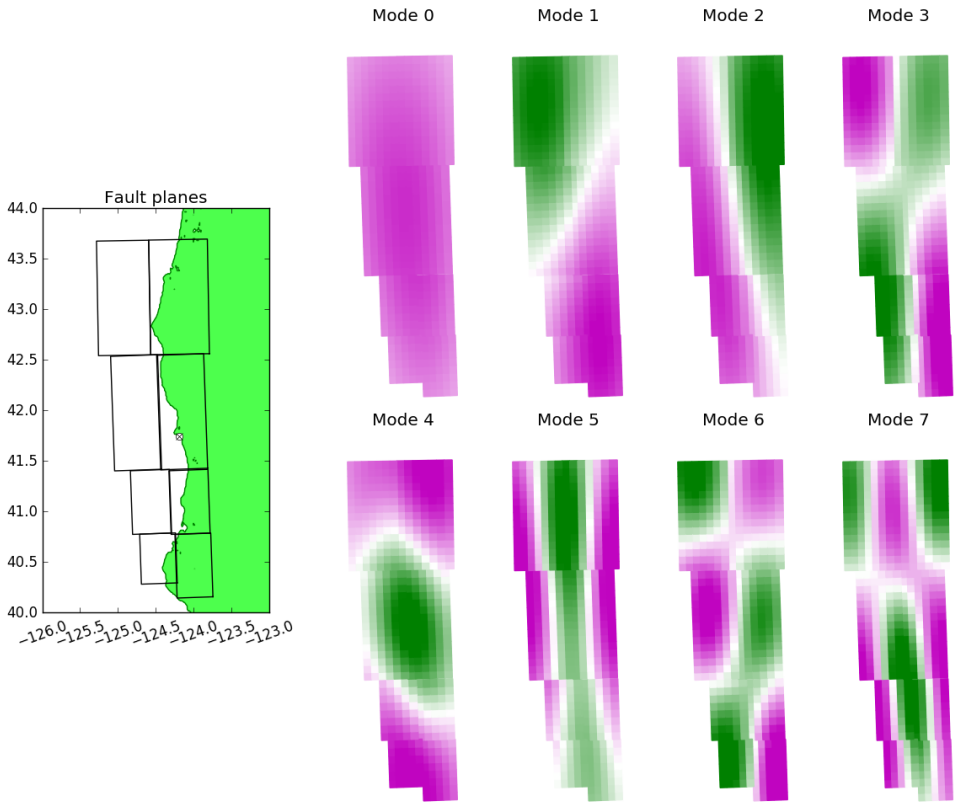

Figure 50: Southern portion of the CSZ fault showing location of Crescent City, CA and the 8 subfaults that are further subdivided into 540 subfaults. The first 7 eigenmodes $\mathbf{v}_j$ of the resulting covariance matrix are also shown.

These are subdivided into 540 smaller fault planes for the purposes of defining the slip. The slip distribution will be represented by a vector of 540 entries, denoted by $\mathbf{s}$. The slip on each subfault was assumed to follow a joint lognormal distribution. To define the 540×540 correlation matrix, we first computed the pairwise distance between subfault i and subfault j. We can compute the Euclidean distance, but for this fault geometry we expect a longer correlation length in the strike direction than down-dip, so we define

$$C_{ij} = \exp(-(d_{strike}(i,j)/r_{strike}) - (d_{dip}(i,j)/r_{dip})) \tag{8}$$

where $d_{strike}(i,j)$ and $d_{dip}(i,j)$ are estimates of the distance between subfaults i and j in the strike and dip direction, respectively, and r_{strike}, r_{dip} are the correlation lengths in each direction. We define $d_{dip}(i,j)$ using the difference in depth between the two subfaults and the dip angle δ as $d_{dip}(i,j) = d_{depth}/\sin(\delta)$, setting $d_{strike}(i,j) = \sqrt{d_{ij}^2 - d_{dip}(i,j)^2}$. We take the correlation lengths to be 40% of the fault length and

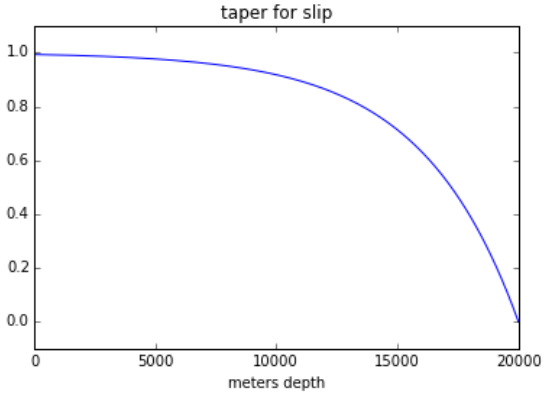

Figure 51: The taper (10) applied to the slip distribution.

width respectively, $r_{strike} = 130$ km and $r_{dip} = 40$ km. Since the lognormal variables are not linear, we perform the K-L expansion on the transformed Gaussian variables, $\log \mathbf{s}$, with corresponding correlation C^g_{ij}. The slip $\mathbf{s}$ is now written as,

$$\mathbf{s} = \exp\left[\sum_{j=0}^{\infty} z_j \sqrt{\lambda_j}\mathbf{v}_j\right] \quad \text{where } z_j \sim \mathcal{N}(0,1), \tag{9}$$

and where λ_j and $\mathbf{v}_j$ are eigenpairs of the covariance matrix for the Gaussian variables. A slip realization is generated by drawing samples from the normal distribution $\mathcal{N}(0,1)$, then computing $\mathbf{s}$ in the series above. In this work, the series was truncated after the term 20. Once a slip is generated from the series, we apply a tapering at the down-dip edge, given by

$$\tau(d) = 1 - \exp(-20(d - d_{max})/d_{max}) \quad \text{with } d_{max} = 20000\text{m}. \tag{10}$$

The plot of this taper is given in Figure 51. We then scale the slip to obtain the desired seismic moment, for example $M_w = 8.6$.

Over the course of this project we experimented with several sets of realizations. In this report we only present results from the set that we found most challenging due to the wide range of magnitudes and flooding behavior in Crescent City.

We generated 100 events of magnitude $M_w = 8.6$ using the K-L expansion with a correlation length of 130 km in the strike direction and 40 km in the dip direction. In order to generate extreme events, these were sampled by first choosing the z_j coefficients each from a Gaussian with standard deviation 1. We draw 20,000 quasi Monte Carlo samples from a Halton sequence, to ensure that the samples are spread out and cover the probability space well. These samples are shown in Figure 52. Then we dilate them by multiplying by 4, with the result that the majority of the events are far out in the tail of the distribution. After exponentiating the Gaussian K-L expansion, this dilation appears in the expression (9) as

$$\mathbf{s} \sim \exp\left[\sum_{j=0}^{20} 4 z_j \sqrt{\lambda_j}\mathbf{v}_j\right] = \left(\exp\left[\sum_{j=0}^{20} z_j \sqrt{\lambda_j}\mathbf{v}_j\right]\right)^4. \tag{11}$$

Taking the power of 4 of the slips above has the effect of generating events that tend to have highly concentrated regions of high slip. That is, the resulting tsunami varies dramatically depending on where the high slip region is located, in part because there can be large values of uplift or subsidence along the coast.

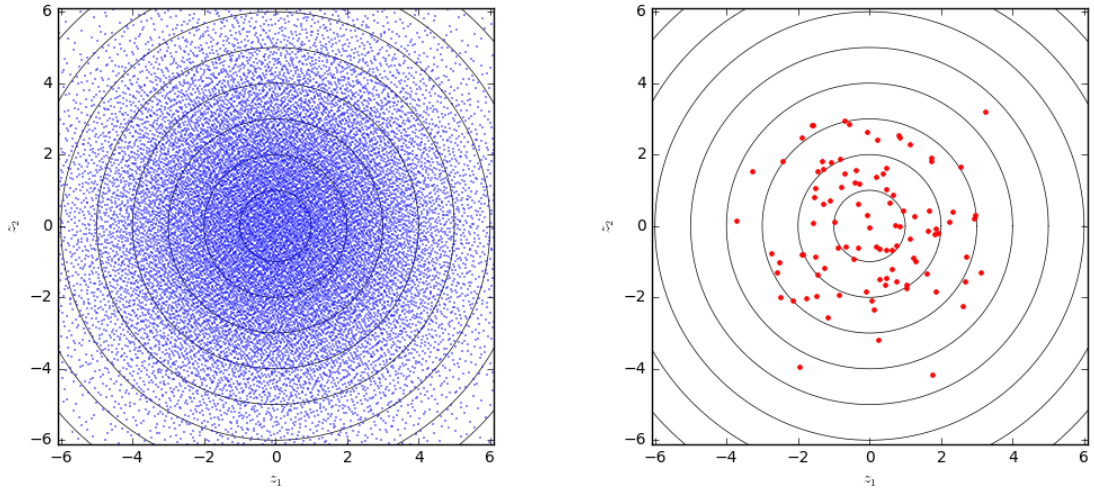

Figure 52: The 20,000 quasi Monte Carlo samples, projected on to z_1-z_2 plane (left), and 100 cluster centers obtained through k-means clustering (right).

Each of these sample points correspond to a potential earthquake scenario, therefore we have so far obtained a catalogue of 20,000 potential earthquake scenarios (corresponding annual probabilities have not yet been assigned.) Since we are constrained by the computational budget, we must reduce the number of scenarios to $N = 100$. To this end, we apply the k-means clustering to the large number of scenarios, and obtain 100 clusters C_k. There are various clustering algorithms that could accomplish this, for example, the well-known Lloyd's algorithm [20]. The cluster centers will each represent an earthquake scenario which will be used for the tsunami simulation. An example of cluster centers is shown in Figure 52.

The remaining task is to assign the annual probability to each of the earthquake scenarios. Relative probabilities were assigned to these 100 events by assuming that each event is representative of some region of parameter space (a high-dimensional space whose dimension is equal to the number of terms used in the K-L expansion, 20 for this work). A Voronoi tessellation of this 20-dimensional space was constructed in performing the k-means clustering above, subdividing the space into disjoint regions R_k, each corresponding to the cluster C_k. This region was formed for each event determined by coefficients $z^{(j)} \in \mathbb{R}^{20}$ and consists of all points z that are closer to $z^{(j)}$ than to any other $z^{(i)}$. Then 1,000,000 Monte Carlo samples of the z values were generated from the joint normal distribution and the fraction of these points falling in each region R^k was used for the probability p_k. These samples will be indexed by z_ℓ. Let us denote the indicator function that returns 1 if a sample z lies in the region by $I_k : \mathbb{R}^{20} \to \{0, 1\}$ for sample z by $I_k(z)$ and the k-th cluster by R_k,

$$I_k(z) = \begin{cases} 1 & \text{if } z \in R_k, \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

Then the probability weights are computed by

$$p_k \equiv \frac{1}{M} \sum_{\ell=1}^{M} I_k(z_\ell) \tag{13}$$

Since the higher-order terms in the K-L expansion do not significantly affect the tsunami resulting from the earthquake, we computed the probability weights based on the 7-dimensional $\mathcal{N}(0, I)$, for computational efficiency.

From this set of 100 realizations we created 300 more by scaling them up to a larger magnitude by multiplying the slip on each subfault (and hence the seafloor deformation) by approximately a factor of 2, which increases the moment magnitude by 0.2. In this manner we generated 100 realizations at each magnitude M_w 8.6, 8.8, 9.0, and 9.2.

The probabilities computed for each of he original 100 realizations using the Voronoi diagram approach described above was then apportioned between the 4 different magnitude events using weights of (0.3, 0.3, 0.3, 0.1), with equal weighting given to the smaller three magnitudes and less weight to the extreme M_w 9.2 events.

We note again that these realizations are **not** meant to be a realistic representation of actual events or probabilities associated with the CSZ. In particular, as noted in Section 3, this approach gave realizations with highly concentrated slip patches where the slip exceeded 180 m for some of the M_w 9.2 events, which is not realistic although it generates interesting test data.

In retrospect, it is also not clear that the approach used here to assign probabilities is correct, even if we were sampling the correct probability density to begin with, due to the high dimension of the space we are attempting to sample. Further work is required to generate a large number of realizations with realistic probabilities attached, which is beyond the scope of the current project.

H Clustered realizations

On the next 4 pages we show at most 5 realizations from each of the 20 clusters generated using the algorithm of Section 5.1. The cluster numbers are arbitrary, but they are arranged in order of increasing η_{mean} values.

The 4 pages after that show at most 5 realizations from each of the 20 clusters generated using the algorithm of Section 5.2. The cluster numbers are arbitrary, and they are **not** arranged in order of increasing η_{mean} values.

H.1 Etamean-Etamax clusters

Cluster 1 (5 realizations)

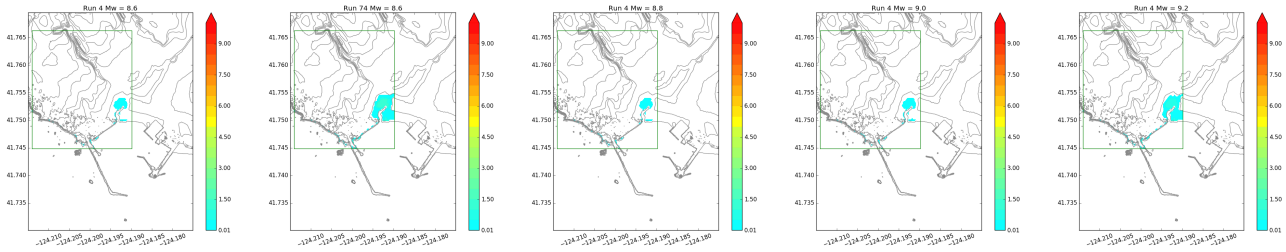

Cluster 2 (41 realizations)

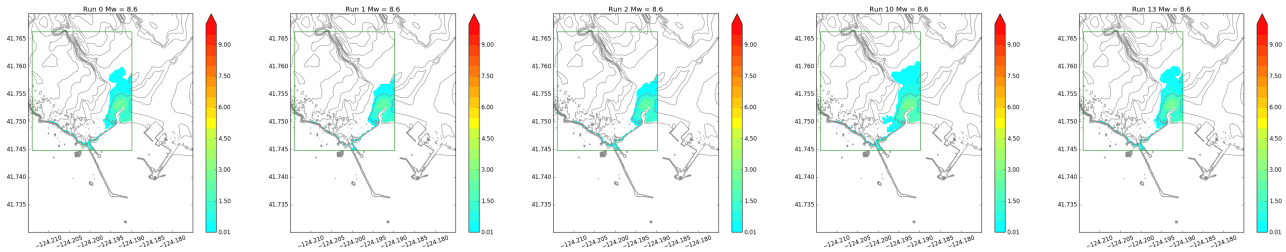

Cluster 3 (49 realizations)

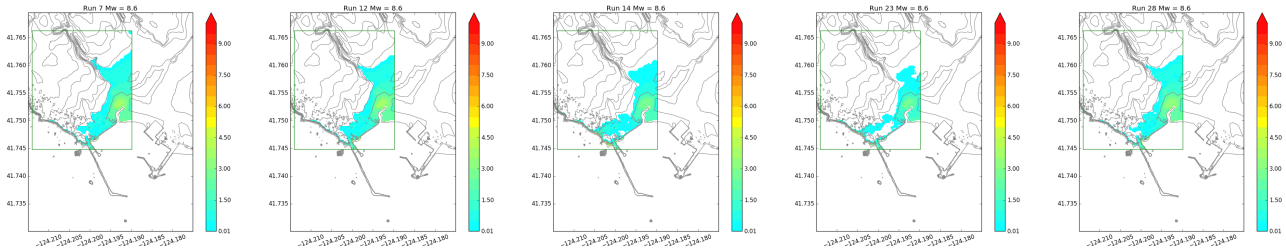

Cluster 4 (51 realizations)

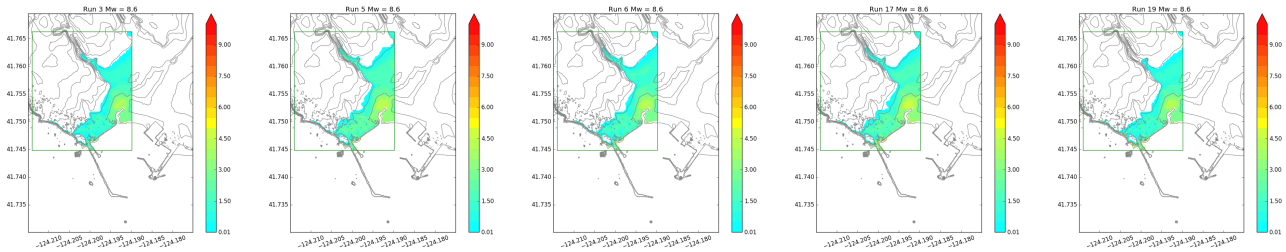

Cluster 5 (50 realizations)

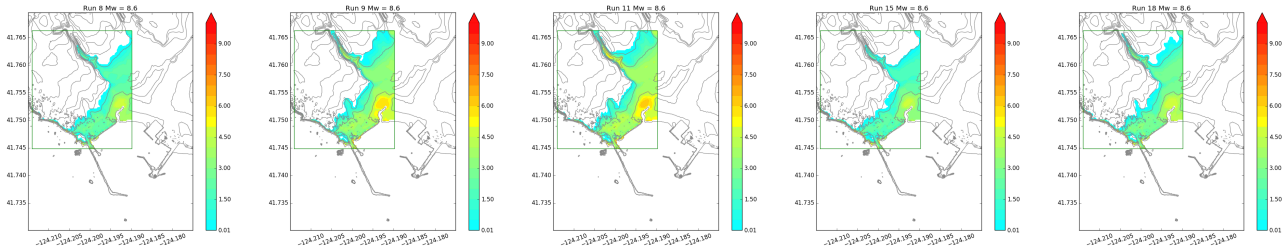

Cluster 6 (16 realizations)

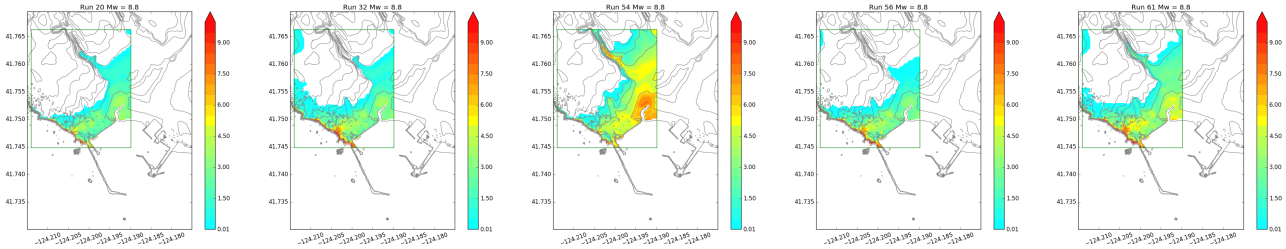

Cluster 7 (23 realizations)

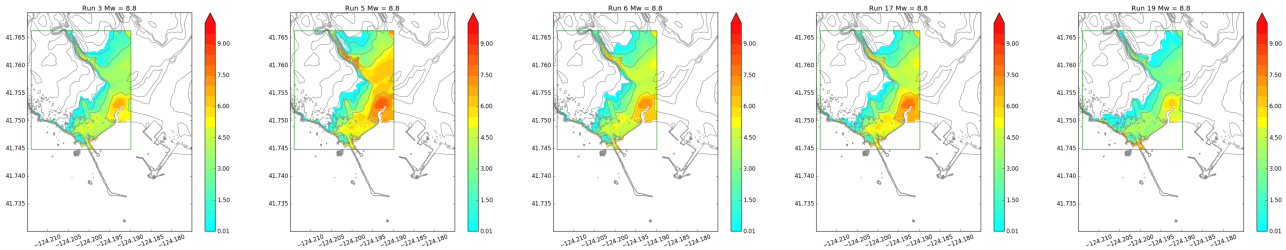

Cluster 8 (18 realizations)

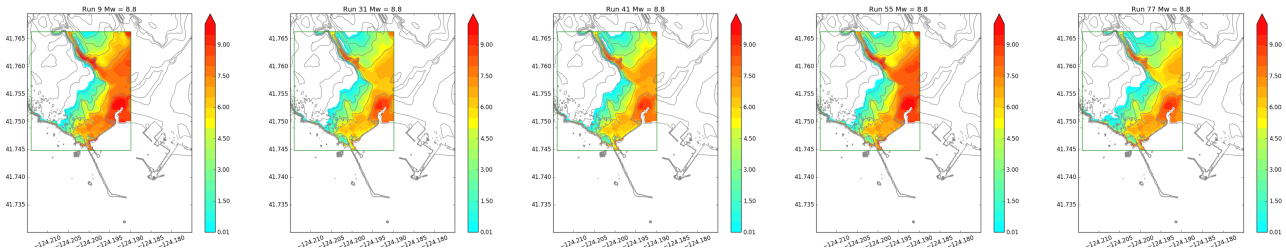

Cluster 9 (11 realizations)

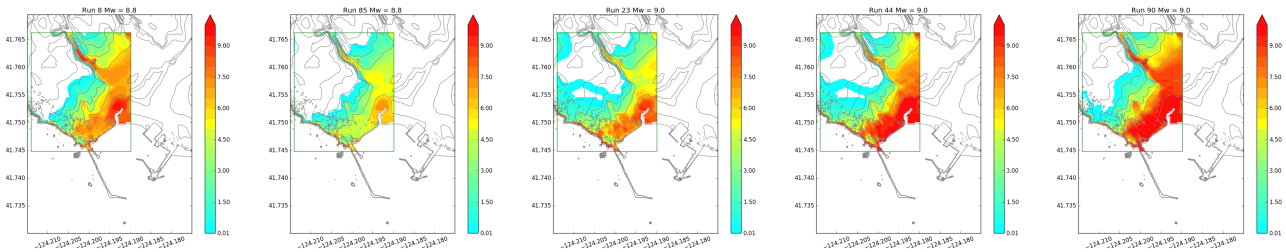

Cluster 10 (24 realizations)

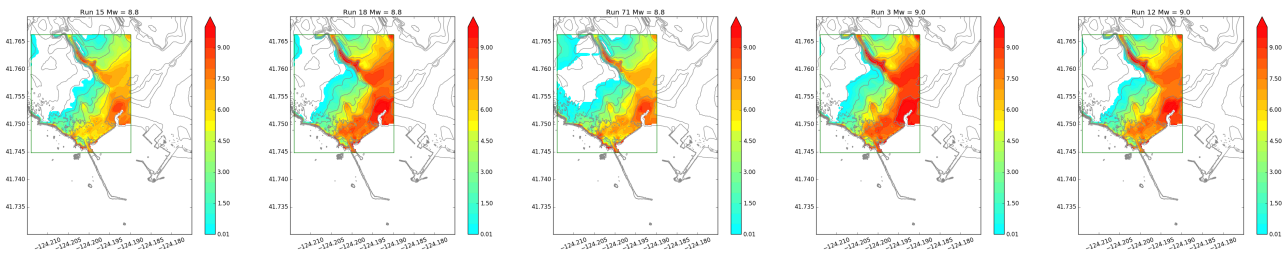

Cluster 11 (18 realizations)

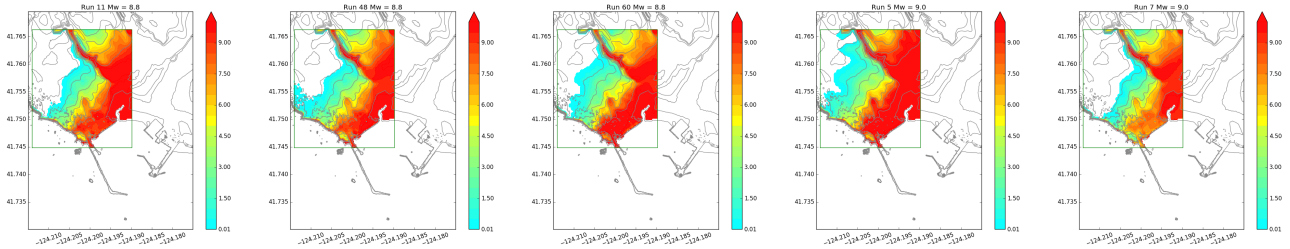

Cluster 12 (17 realizations)

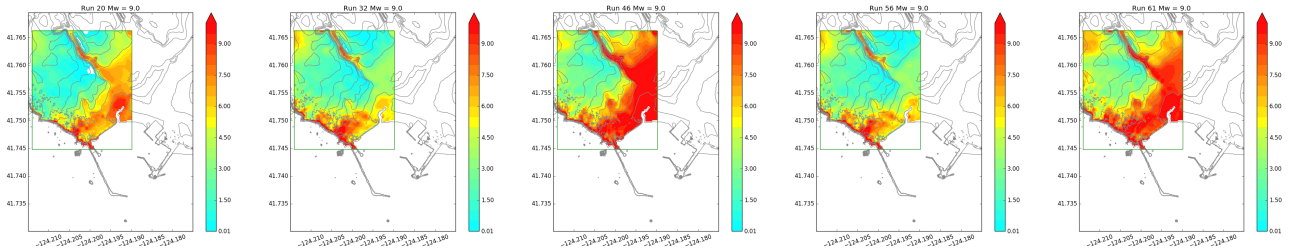

Cluster 13 (16 realizations)

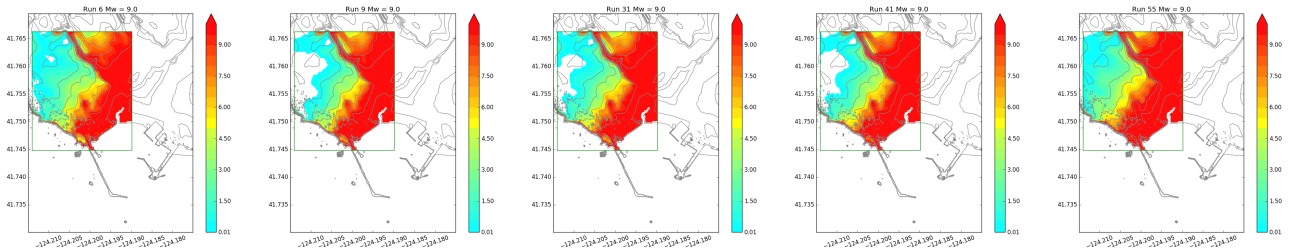

Cluster 14 (5 realizations)

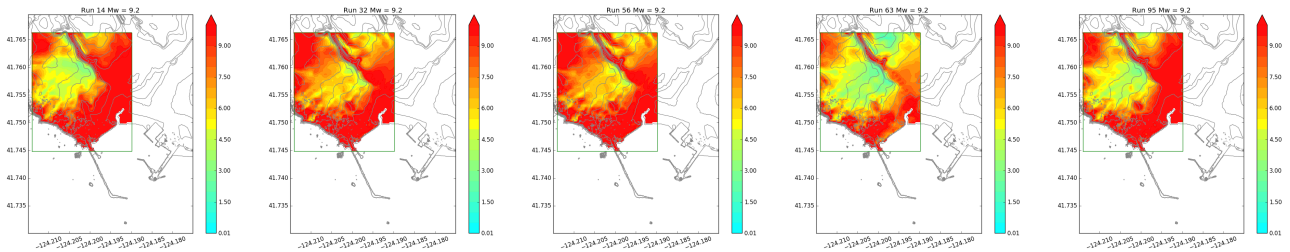

Cluster 15 (16 realizations)

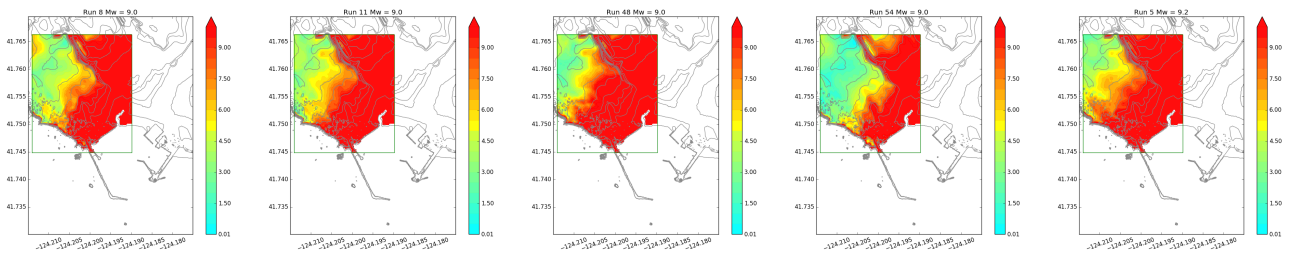

Cluster 16 (15 realizations)

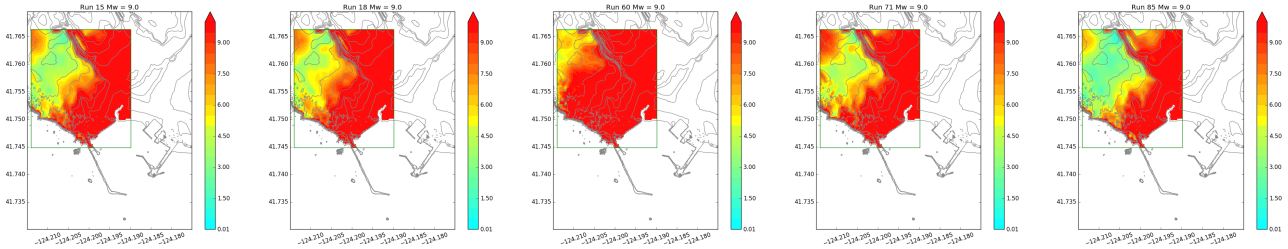

Cluster 17 (7 realizations)

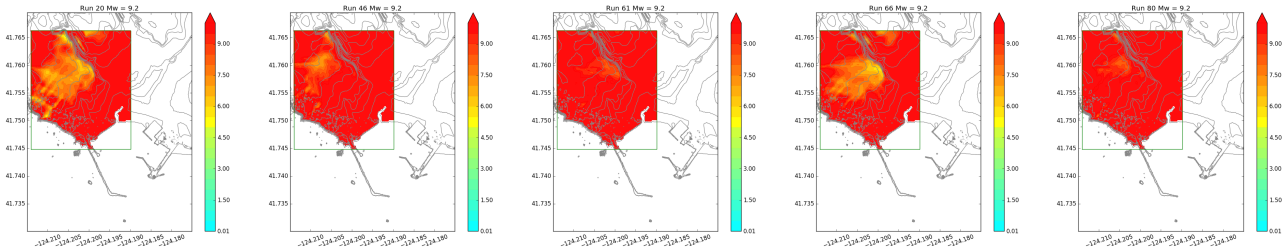

Cluster 18 (7 realizations)

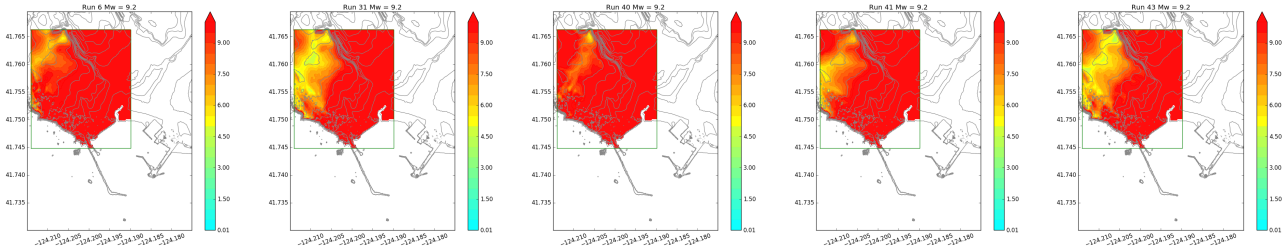

Cluster 19 (4 realizations)

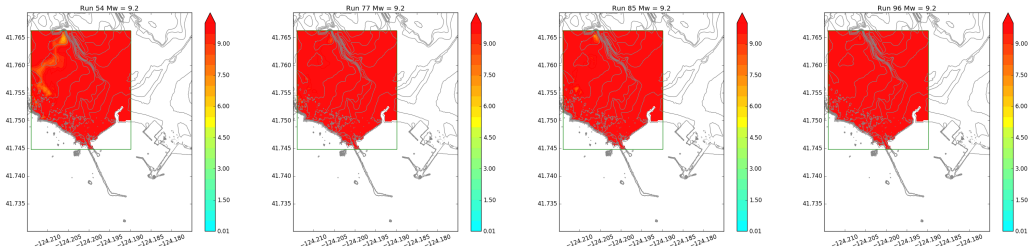

Cluster 20 (7 realizations)

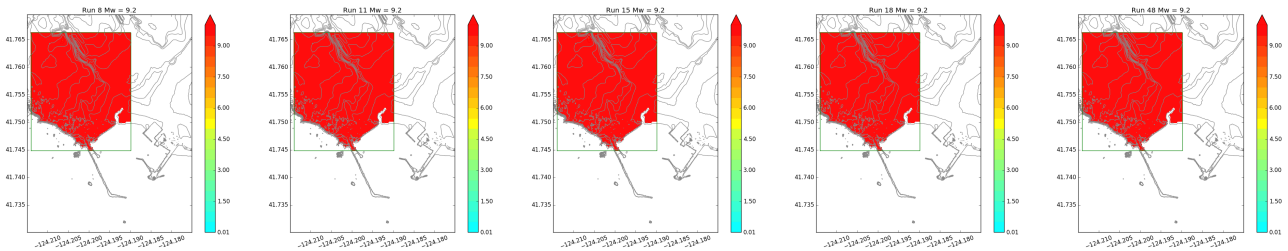

H.2 dtopo clusters

Cluster 1 (6 realizations)

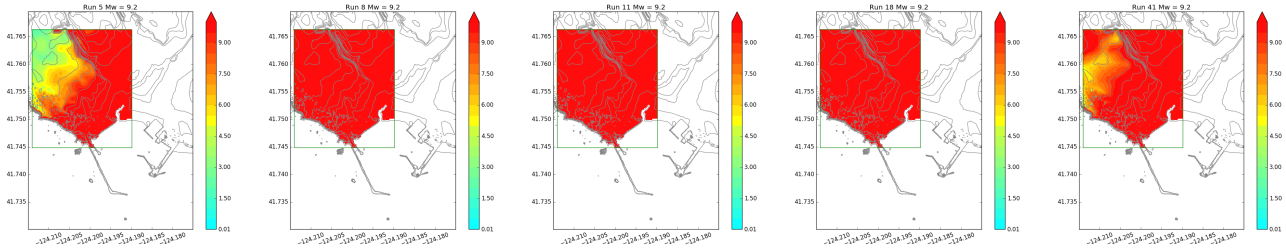

Cluster 2 (3 realizations)

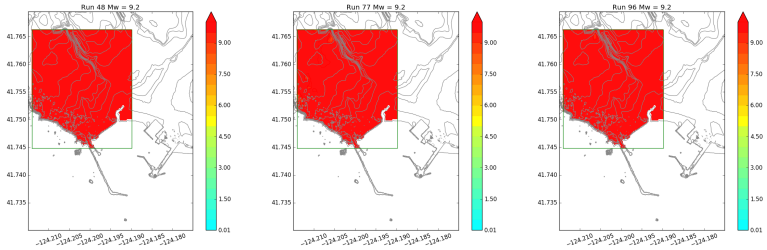

Cluster 3 (8 realizations)

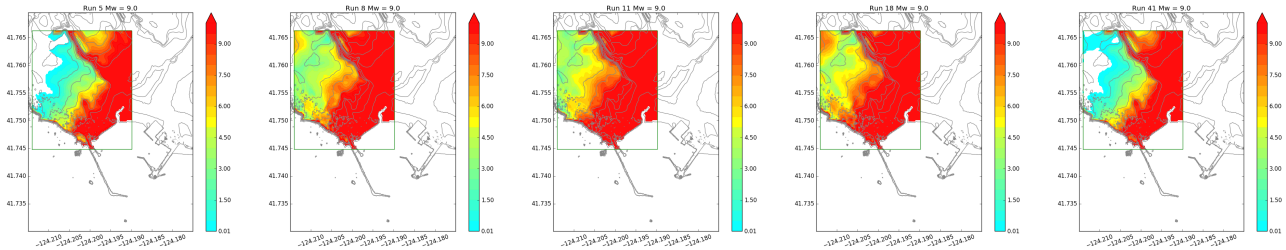

Cluster 4 (15 realizations)

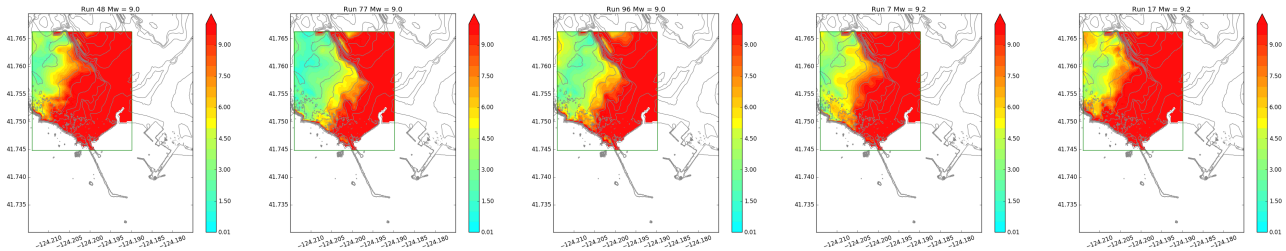

Cluster 5 (2 realizations)

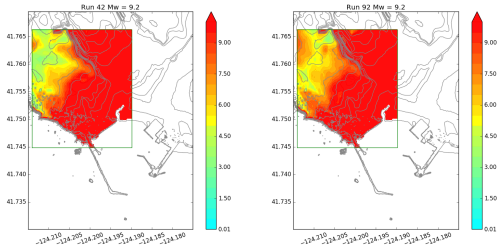

Cluster 6 (15 realizations)

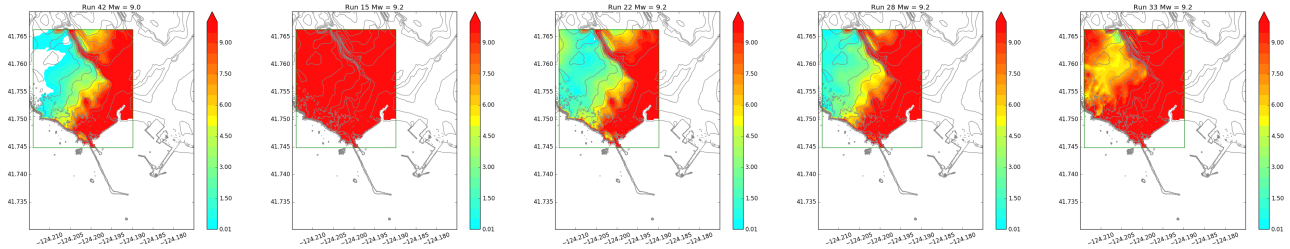

Cluster 7 (22 realizations)

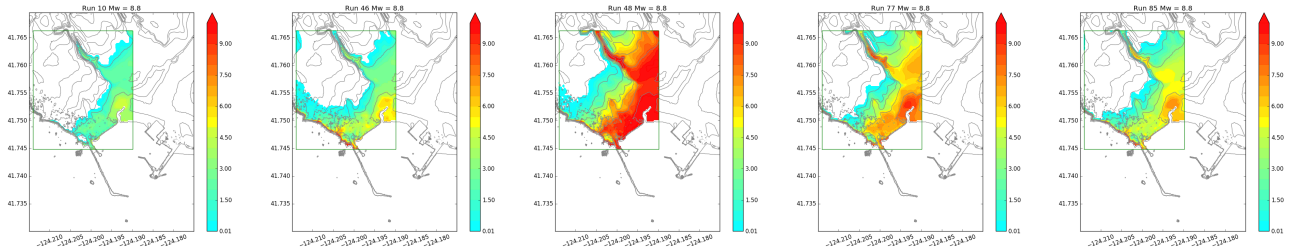

Cluster 8 (10 realizations)

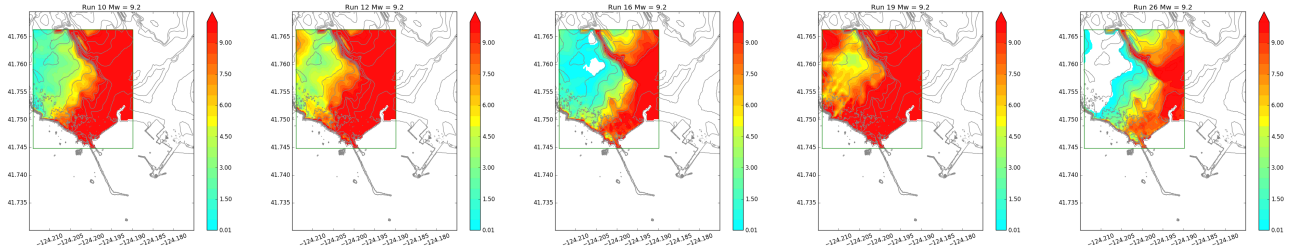

Cluster 9 (24 realizations)

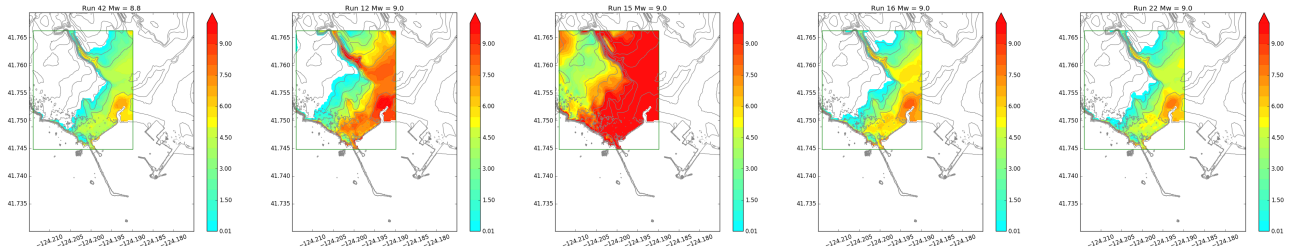

Cluster 10 (49 realizations)

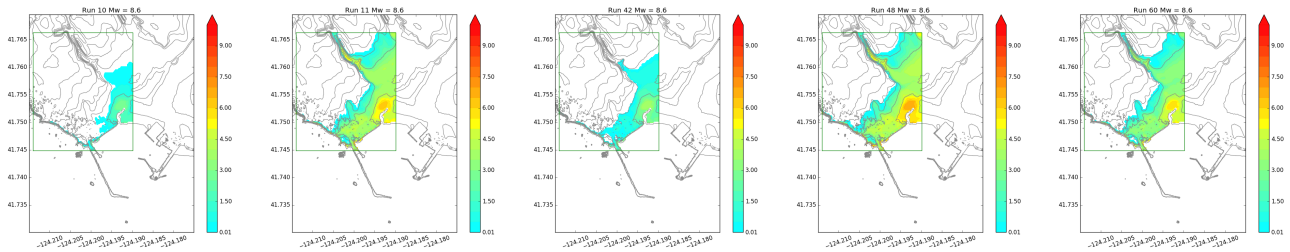

Cluster 11 (16 realizations)

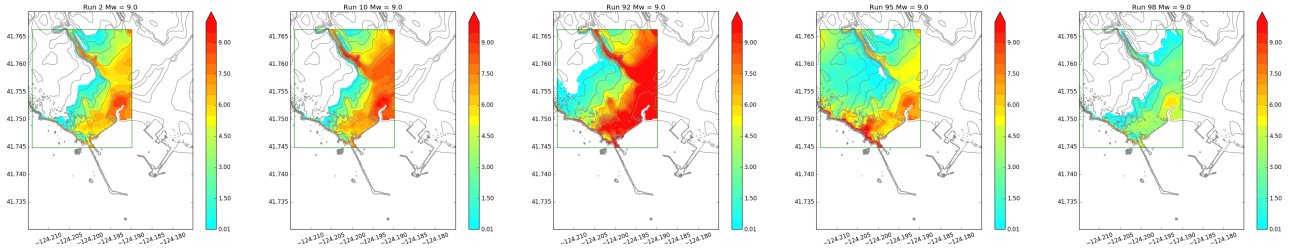

Cluster 12 (106 realizations)

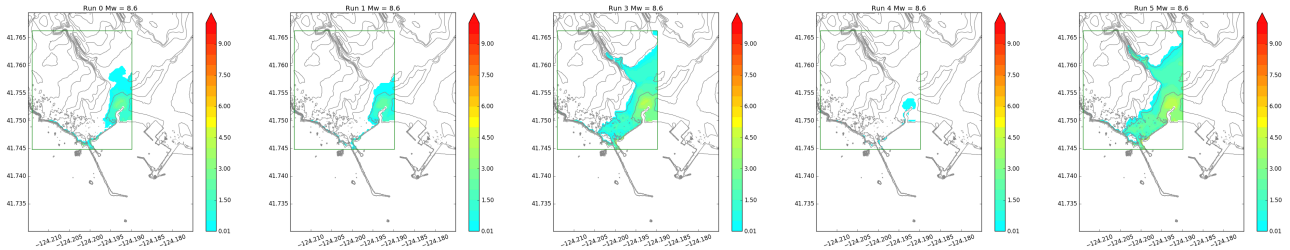

Cluster 13 (50 realizations)

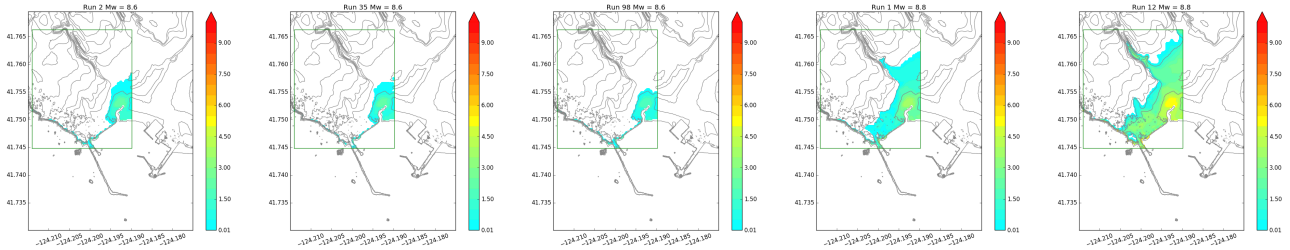

Cluster 14 (37 realizations)

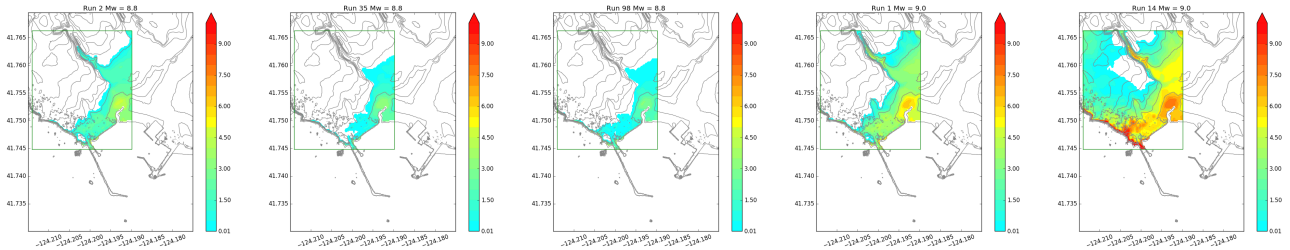

Cluster 15 (17 realizations)

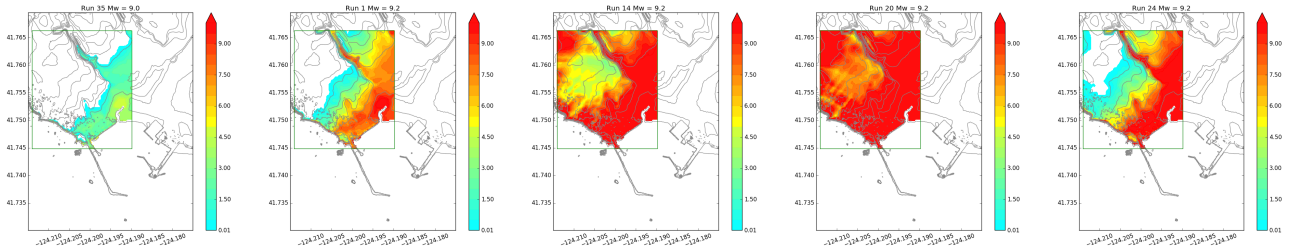

Cluster 16 (3 realizations)

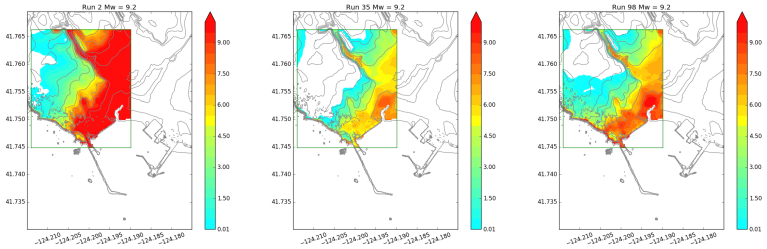

Cluster 17 (4 realizations)

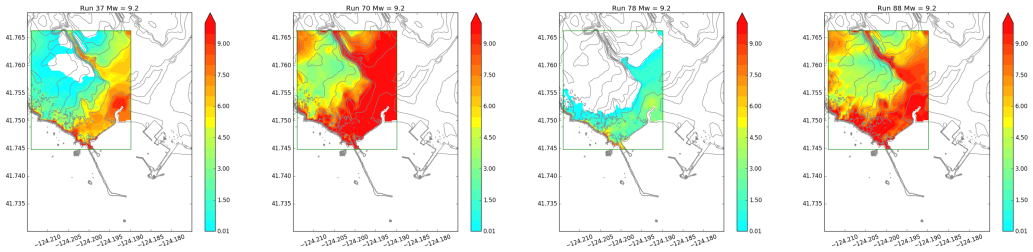

Cluster 18 (8 realizations)

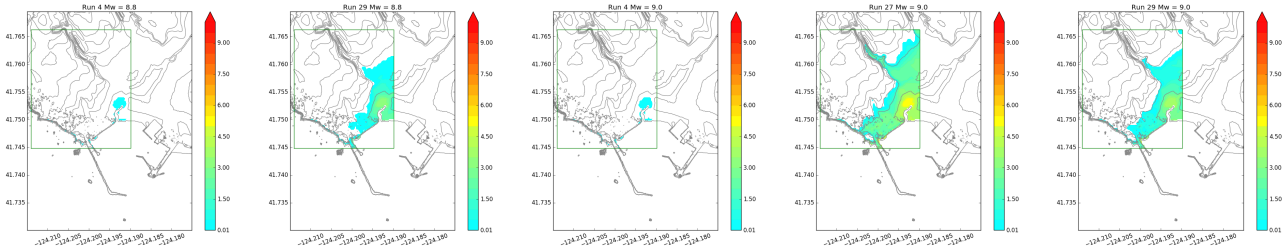

Cluster 19 (3 realizations)

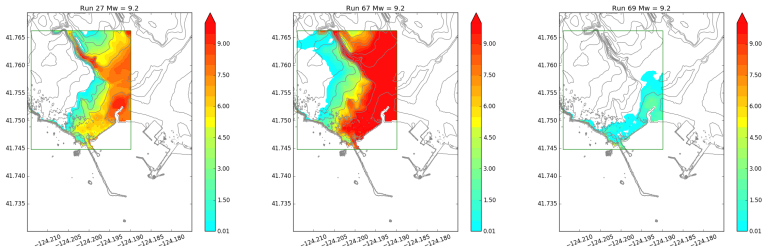

Cluster 20 (2 realizations)