The Annals of Applied Statistics
2018, Vol. 12, No. 4, 2023–2053
https://doi.org/10.1214/18-AOAS1142

FUNCTIONAL EMULATION OF HIGH RESOLUTION TSUNAMI MODELLING OVER CASCADIA

BY SERGE GUILLAS[*,1], ANDRIA SARRI[*], SIMON J. DAY[*,1], XIAOYU LIU[*] AND FREDERIC DIAS[†]

University College London[*] and University College Dublin[†]

The rarity of tsunamis impels the scientific community to rely on numerical simulation for planning and risk assessment purposes because of the low availability of actual data from historic events. Numerical models, also called simulators, typically produce time series of outputs. Due to the large computational cost of such simulators, statistical emulation is required to carry out uncertainty quantification tasks, as emulators efficiently approximate simulators. There is thus a need to create emulators that respect the nature of time series outputs. We introduce here a novel statistical emulation of the input-output dependence of these computer models. We employ the Outer Product Emulator with two enhancements. Functional registration and Functional Principal Components techniques improve the predictions of the emulator. Our phase registration method captures fine variations in amplitude. Smoothness in the time series of outputs is modelled, and we are thus able to select more representative, and more parsimonious, regression functions than a fixed basis method such as a Fourier basis. We apply this approach to the high resolution tsunami wave propagation and coastal inundation for the Cascadia region in the Pacific Northwest. The coseismic representation in this analysis is novel, and more realistic than in previous studies. With the help of the emulator, we can carry out sensitivity analysis of the maximum wave elevation with respect to the source characteristics, and we are able to propagate uncertainties from the source characteristics to wave heights in order to issue probabilistic statements about tsunami hazard for Cascadia.

1. Introduction. Tsunami waves are mostly generated by seabed displacements such as underwater earthquakes and submarine landslides. After generation, the tsunami wave propagates rapidly over the deep ocean, gets amplified when the water depth decreases, and runs up the coast, provoking potentially severe damage. Recent tsunamis have been particularly devastating, as their characteristics were somehow unexpected due to a lack of adequate probabilistic investigation of the source mechanism. Notably, the 2004 Indian Ocean tsunami killed roughly

Received March 2017; revised January 2018.

[1]Supported through the NERC grants PURE (Probability, Uncertainty and Risk in the Natural Environment) NE/J017434/1, and "A demonstration tsunami catastrophe risk model for the insurance industry" NE/L002752/1.

Key words and phrases. Uncertainty quantification, emulation, tsunami modelling, sensitivity analysis, functional data analysis.

225,000 people across the Indian Ocean, especially from vulnerable communities, and the 2011 Tohoku tsunami killed nearly 20,000 people in Japan. The key question is thus how to quantify the possible hazard, and its impact on population and infrastructure, thereby gaining wisdom in the tsunami risk assessment [Okal (2015)]. This question needs to be examined through a better scientific understanding of the source, propagation, and inundation. However, data from historic events are scarce. Therefore, advanced numerical simulations have been used over the last 30 years for planning and risk assessment purposes. These tsunami simulators are very costly to run at high resolution. Hence, an attractive approach is to build and run statistical surrogates, also called emulators, that approximate the numerical simulators.

1.1. *Emulation of high-dimensional outputs*. Many complex computer models, such as engineering, climate or fluid dynamics models, produce outputs that are typically time series at various locations in the physical output space. It is advantageous to replace the simulator by an emulator for various uncertainty quantification tasks. Indeed, Monte Carlo approaches typically require thousands of runs of the simulator and are thus impractical for high accuracy models that are computationally expensive. The intrinsically high-dimensional nature of the outputs (especially for high frequency time series) increases dramatically the complexity of building an adequate multivariate emulator. The simplest approach would be to build separate independent emulators for each output. However, this method has major drawbacks. It inflates computational costs (though in a linear manner), but more importantly ignores the correlations between the outputs across time in the process of building the emulator (though the emulated outputs may turn out be correlated); it lacks the ability to inherently borrow strength across outputs.

Efficient emulations of simulators with multiple outputs have been proposed [Rougier (2008), Bilionis et al. (2013)]. Rougier (2008) introduced the Outer Product Emulator (OPE) that creates one emulator for all outputs. It simplifies the representation of the output functions by using products of functions with respective arguments, input parameters and time. This means that there is no simultaneous influence of inputs and time on outputs beyond a multiplicative effect of separate input and time functions. As a result, the cost of building and running the OPE emulator is much smaller compared to a general multivariate emulator. The OPE has the form

$$f_i(r) = \sum_{j=1}^{\nu} \beta_j g_j(r, s_i) + \varepsilon(r, s_i), \tag{1}$$

where $f_i(r)$ is the ith simulator output at input r, the g_j are the regressor functions, the β_j are the unknown coefficients and ε is the residual assumed to be a Gaussian Process (GP), s_i represents the output domain—for example, time or space—corresponding to the ith simulator run. In order to build an emulator, appropriate prior distributions for β and ε must be chosen. A convenient choice is

the Normal Inverse Gamma (NIG) framework that enables the use of conjugacy, described by

$$\beta|\tau, B \sim N(m, \tau V), \qquad \varepsilon|\tau, B \sim \mathrm{GP}\big(0, \tau\kappa_\lambda(\cdot)\big), \qquad \tau|B \sim \mathrm{IG}(a, d), \tag{2}$$

where $B = \{m, V, a, d, \kappa_\lambda(\cdot)\}$ is the set of the hyperparameters, m and V are respectively a mean vector and a covariance matrix of a multivariate Normal distribution for the coefficients β in the regression, τ is a variance inflation factor, $\kappa_\lambda(\cdot)$ is the covariance function of the residuals with correlation lengths λ and IG denotes the Inverse Gamma distribution. Summing up, $\{\beta, \varepsilon\} \sim \mathrm{NIG}(m, V, a, d)$ where the hyperparameters a and d denote the degrees of freedom and the scale, respectively.

Two main characteristics distinguish the OPE from a standard multivariate emulator. The first is that the covariance function of the residuals is separated into inputs r and outputs s, that is,

$$\kappa_\lambda\big(r, s, r', s'\big) = \kappa_\lambda^r\big(r, r'\big) \times \kappa_\lambda^s\big(s, s'\big). \tag{3}$$

The second is that the set of the regressor functions, G, is the outer product of the set of regressors for inputs, $G^r \overset{\text{def}}{=} \{g_j^r(r)\}_{j=1}^{j=\nu_r}$, with the set of regressors for outputs, $G^s \overset{\text{def}}{=} \{g_j^s(s)\}_{j=1}^{j=\nu_s}$, or in other words, the functions g_j are given by products $g_j(r, s) = g_j^r(r) \otimes g_j^s(s)$ where $\otimes$ is the outer product symbol and $j = \{1, \ldots, \nu\}$, where $\nu = \nu_r \times \nu_s$. Appropriate choices for the regressor functions are necessary.

Sarri, Guillas and Dias (2012) used the OPE for an idealized tsunami modelling application, where a simplified landslide generated the wave and an analytical solution was used to compute the wave elevation. Without the emulation of time series outputs, emulation was also recently used to propagate uncertainties for a real landslide event 21 ka ago in the Northeast Atlantic [Salmanidou et al. (2017)], and for the purpose of quantifying the influence of uncertain bathymetry near shore [Liu and Guillas (2017)]. The present paper provides a significant extension to a realistic tsunami modelling application, where the tsunami is generated by a truthful seabed deformation linked to an underwater earthquake and the tsunami propagation and amplification are obtained by solving numerically the nonlinear shallow water equations over a complex and large domain. We employ the tsunami model VOLNA [Dutykh, Poncet and Dias (2011)]. Its numerical solution uses an unstructured mesh in the fluid domain combined with high resolution bathymetry and topography. In the present paper, we also use the OPE. However, allowing for a variation in the duration of the earthquake necessitates an alignment of the various time series of outputs, that we perform through functional registration. Finally, the shapes of the waves are now irregular due to the influence of the realistic bathymetry on the propagation, so the regression functions are chosen to be adaptive to the time series of outputs. To do so, we employ a data-driven method for the representation of the outputs, namely the Functional Principal Components Analysis (FPCA). As a result, we demonstrate that a novel type of Probabilistic

Tsunami Hazard Assessment can be potentially achieved for tsunami hazard, as the full time history can be accounted for, without sacrificing numerical precision and quantification of uncertainties.

1.2. *Using Functional Principal Components in emulation*. In the present paper, a novel approach for building a GP emulator when outputs are curves is introduced. Functional Principal Components (FPCs) [Ramsay and Silverman (2005)] are used in place of the standard choice of either fixed basis functions (e.g., Fourier basis, wavelets, polynomials) or adaptive multivariate PCs, for the outputs representation [Bayarri et al. (2007), Higdon et al. (2008)]. In these previous approaches, the coefficients, or weights, associated with these basis functions are individually modelled as GPs. In our approach, we fit constant coefficients in a regression onto these basis functions, and what is left over is modelled as a separable GP. FPCs are similar to the standard multivariate PCs except that time series are viewed as curves, not just vectors with no order nor continuous features. The smoothness is acknowledged and employed as a constraint in the modelling so that the resulting FPCs are smooth (and thus physically more meaningful); this approach mitigates the curse of dimensionality whenever a few underlying (necessarily smooth due to the physics) modes of variability are inferred.

Multivariate PCs, also called Empirical Orthogonal Functions in geophysical sciences, have been used in the past to reduce dimension [Higdon et al. (2008), Kleiber, Katz and Rajagopalan (2013), Chang et al. (2014)] in GP emulation (and its use in calibration). We propose here to make use of the smoothness of the time series of wave elevations to define our smooth basis functions as FPCs, with clear benefits, as shown below. We note that three alternative methods to our approach have been published. One was developed by Hung, Joseph and Melkote (2015) who emulated simulators with functional outputs, without the use of data-driven FPCs like we do, but with a marginal profile approach in the input and functional output domains to identify the functional form of the mean. The second is using a Bayesian forward selection procedure implemented in blind kriging [Joseph, Hung and Sudjianto (2008)]: the mean is chosen from an initial set of functions. The third [Morris (2012)] models functional inputs and outputs within the correlation function itself by generalizing the sums in product correlations to integrals.

Wavelets typically fit better curves whose variations are sharper than Fourier basis functions. However, the adaptive nature of FPCs allows for complex modes of variations that cannot be accommodated by Fourier modes or wavelets (or combinations of the two), unless using a prohibitive number of basis functions; thus FPCs allow a wider range of adequate representations. Finally, the cost of building a GP emulator is cubic in terms of data points, since the Cholesky decomposition is employed to compute the likelihood. However, for the OPE, the complexity is proportional to the sum of a cubic number of inputs and a cubic number of time points, thanks to the separability of the GP. This is particularly helpful when a large number of emulators needs to be constructed; for instance, this would help

for hazard planning as in Spiller et al., 2014. This also might help when time is a constraint. Early warning systems could use emulators: the current approach of looking up tables of pre-computed runs cannot produce a proper uncertainty quantification of the warning. Note that methods that put GP weights on a few basis functions (e.g., FPCs or Fourier) will be computationally faster than the OPE as only a few coefficients are fitted as GPs (the complexity in time disappears). However, such methods force the interaction between input and time to be solely along these directions, whereas our method still allows for a GP over inputs and time to describe, in a completely flexible manner, the residual variations not captured by the main modes of variability.

1.3. *Tsunami hazard*: *Modelling and emulation*. Our application to long-term probabilistic tsunami hazard assessment aims to reduce the risks of loss of life and building and infrastructure damage from a potentially tsunamigenic earthquake.[2] A particular problem for the 2011 Tohoku tsunami was that uncertainties in the tsunami hazard assessments were underestimated, leading to adoption of mitigation measures that worked well up to their design hazard intensity limits, but failed with disastrous consequences in the face of an event larger than anticipated [Okal (2015), Day and Fearnley (2015)]. Our proposed method requires the formulation and acknowledgement of the uncertainties in the source, so the analysis is more complete, and thus helpful to use in tsunami hazard assessments.

Tsunami hazard assessments are heavily dependent upon the use of simulators, because tsunamis are relatively rare but highly destructive events: data on past tsunamis are sparse in most regions. To ensure accurate modeling of tsunami wave physics and provide sufficiently high resolution variations in the hazard due to local topography and bathymetry, tsunami simulators are computationally expensive. Moreover, the balance between full representation of tsunami wave physics and model stability is difficult to maintain [Behrens and Dias (2015), Bernard and Titov (2015)].

The calculation of the uncertainties in predictions is challenging due to the large number of model evaluations required. We show here how emulators can be used for such time—and resource—demanding analyses. However, Spiller et al. (2014), who studied an application similar to ours (volcanic flow hazard), concluded that a joint space-time modelling, like in for example, Kleiber, Katz and Rajagopalan (2013), should not be carried out due to the very local nature of such hazards. We agree with Spiller et al. (2014) that spatial outputs must be treated individually as site-specific since outputs (here water flows) are heavily dependent on the surrounding bathymetry and topography. There are also uncertainties in the bathymetry (less so in the topography) that could be accounted for [Liu and Guillas (2017)].

[2]Strictly speaking a tsunamigenic fault rupture. Indeed, both the seismic shaking and the tsunami are ultimately consequences of the rapid slip that occurs in the fault rupture; the seismic shaking is not the direct cause of the tsunami.

1.4. *Outline*. In the present paper, we emulate entire time series of wave elevation outputs as a function of three major inputs: rupture velocity of the earthquake, maximum uplift of the seabed overall and at the trench (where the seabed deformation can be discontinuous). Indeed, the damage due to the impact of waves on structures relates to the shape of the train of waves through repeated impacts and scouring. However, for our sensitivity analysis and uncertainty propagation study, we only picked the maximum wave elevation to merely illustrate how the uncertainty analysis can work on a simple output for complex tsunami source. In the future, we could carry out a sensitivity analysis and a uncertainty propagation study for the entire history since we emulate it; this would present challenges in the representations of the resulting uncertainties across time (but provide little value to the current tsunami engineering models yet unable to fully account for time history) and thus is beyond the scope of this paper. The tsunami engineering field is only nascent, but any ambitious risk model for structures on the coast should include such time series information (ideally jointly with flow velocities). Obviously, proxies of outputs could be valuable for tsunami engineering [Bricker and Nakayama (2014)], for example, integrated heights over time (a proxy for accumulated pressure loading on structures), sum of absolute differences in maxima and minima (a proxy for flow intensity).

Our work uses the case of Cascadia, the region covering the Pacific North West of Canada and USA, because of particular challenges of tsunami hazard assessment there, but our emulator approach has wide application to other regions. There is a pressing need to quantify such risk in detail since Leonard, Rogers and Mazzotti (2014) estimate the probability of run-up with significant damage potential (> 3 m) in Cascadia to be 10–30% in the next 50 years. In Section 2, tsunamigenic seabed uplifts for Cascadia are presented, followed by Section 3 where the uncertainties in these uplifts are defined. In Section 4, the functional landmark registration and Functional Principal Component Analysis (FPCA) are applied to our tsunami model outputs. Finally, Section 5 describes the functional emulation and its application to sensitivity and uncertainty analyses, before the conclusion.

2. Realistic representation of coseismic seabed deformation for Cascadia. The representation of the source characteristics and displacement and its numerical modeling rely on an improved comprehension of the fault rupture mechanism and the coseismic sea bottom displacement that constitutes the source of fault rupture-generated tsunamis, as well as high performance computing for tsunami modeling [Behrens and Dias (2015)]. Here, we emphasise that the design of the source representation must take account of the epistemic uncertainties in the understanding of tsunami sources in that region, as these differ between regions.

The Cascadia Subduction Zone (CSZ) (see Figure 1) experienced its most recent locally-generated tsunami on 26 January 1700 [Satake et al. (1996)], as a result of the most recent giant earthquake on the CSZ. Geological evidence [Leonard et al. (2010), Hawkes et al. (2011), Goldfinger et al. (2012), Wang et al. (2013)]

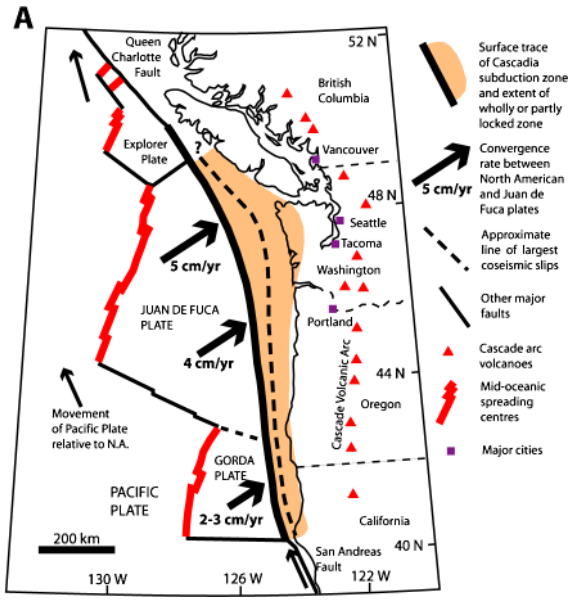

FIG. 1. *Cascadia Subduction Zone* (*CSZ*): *geometry and plate convergence rates.*

indicates that this had a rupture length of about 1000 km—the entire length of the Cascadia subduction zone—and an inferred earthquake moment magnitude of about 9.0, comparable to the Indian Ocean 2004 earthquake. Historical evidence for the earthquake itself is limited to oral traditions from the region [Ludwin et al. (2005)]; the date and time of the earthquake are instead inferred from records of the arrival of the tsunami in Japan [Satake et al. (1996), Satake, Wang and Atwater (2003)]. Seismic activity on the Cascadia subduction zone itself during the time period of instrumental seismic records, and indeed since the establishment of European settlements in the region during the late 18th and early 19th centuries, has been very limited in comparison to most other subduction zones. Nevertheless, abundant evidence for the occurrence of giant earthquakes and the resulting tsunamis in Cascadia has been generated by a wide range of geological and geophysical studies as discussed below.

Given the 3 to 5 cm/year plate convergence accommodated [see. e.g., Mazzotti et al. (2003)], such giant earthquakes would be expected to occur in Cascadia at an average rate of one per 400–600 years. This has proved to be the case, with around 20 giant earthquakes and 20 slightly smaller earthquakes rupturing the Cascadia subduction zone during the past 10,000 years or so [Goldfinger et al. (2012)]. Importantly, the evidence collected to date [e.g., Adams (1990), Wang et al. (1995), Leonard et al. (2010), Nelson et al. (2008), Goldfinger et al. (2012)] provides only an incomplete and indirect insight regarding the earthquake moment magnitudes and other key seismic parameters of the tsunamigenic fault ruptures, but more

complete and direct evidence regarding the coseismic seabed and coastal displacements and the tsunamis themselves. Indeed, much of the evidence regarding the source fault ruptures and resulting earthquake shaking is model-based evidence derived from the interpretation of the geological proxies for surface displacements and tsunamis in terms of simple models relating the reconstructed surface displacements to the source fault rupture, such as the Okada model [Okada (1992)]. Here we specify the tsunami sources in terms of a simple representation of their seabed surface deformation characteristics. This representation is highly simplified compared to the complexity of seabed deformation evidenced in hindcast modeling of recent well-studied tsunamis such as Tohoku 2011 [Simons et al. (2011)]. In particular, it does not include an explicit representation of coseismic movements on splay faults [Moore et al. (2007), Heidarzadeh, Pirooz and Zaker (2009), Witter et al. (2013)] that, while making a limited contribution to the overall seismic moment magnitude of the earthquake, can lead to large uplifts of parts of the seabed above the subduction zone and thus to large tsunami waves as discussed below. We explain how splay faults are accommodated in our simplified representation below. We emphasise that our emulation is computationally efficient and can therefore be extended to more complex representations of tsunami sources with more source characteristics being defined (e.g., a more variable location of the seabed deformation). However, this paper is the first to propose an emulation of the relationship between source characteristics of a realistic earthquake-generated tsunami and tsunami wave heights, and thus our aim is to build on this innovation in future studies with more complex descriptions of the source.

Using publicly available digital elevation models (DEMs) for the initial bathymetry of the ocean floor and coastal shelf seabed, from the National Oceanic and Atmospheric Administration and National Geophysical Data Centre (NOAA/NGDC) website http://www.ngdc.noaa.gov/, different hypothetical scenarios of seabed deformation events that can result in a tsunami are generated in this study by changing three of the source characteristics (see Figure 2): the overall source maximum uplift $z_{\max}$ chosen to vary between 1 and 5 and multiplied by a location-dependent inflation factor of up to 2.2 (in the middle of the margin) [see the North–South and West–East sections in panels (b) and (c)], the ratio $z_t/z_{\max}$ between the uplift coefficient at the trench z_t and the maximum uplift coefficient $z_{\max}$ chosen to vary between 0 and 1 throughout, and the speed of propagation of the source deformation (mimicking the propagation of the fault rupture along the underling fault surface). Note that there is an uplift in the West, and subsidence in the East: Figure 2(b) displays the deformations resulting from the lowest, medium and largest values of $z_{\max}$ and $z_t/z_{\max}$. Different combinations of these characteristics describe different event cases; see movies in the Supplementary Material [Guillas et al. (2018)] and one such final seabed deformation in Figure 3.

We use the inflation factor $z_{\max}$ to allow for large, localised uplifts associated with splay faults as proposed for the central part of the Cascadia subduction zone (corresponding to areas offshore northern Oregon and Washington states) by

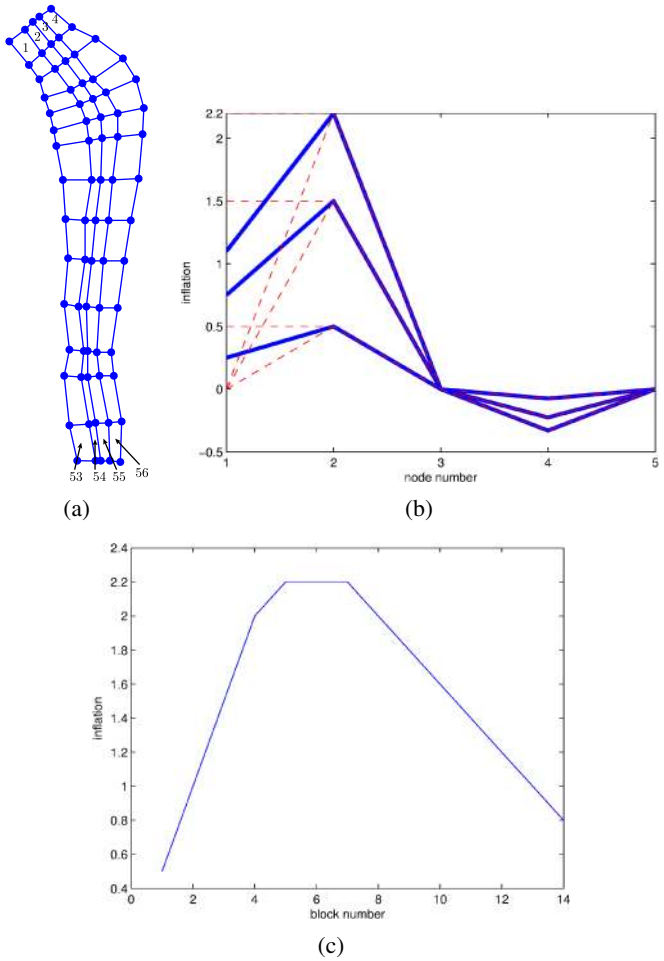

FIG. 2. (a) *Map* (*top*: *North*) *of the deformation area modelled by* 4 *columns* (*defined by* 5 *nodes*) *and* 14 *rows of polygons* (56 *in total*) *matching the fault geometry.* (b) *West–East sections of the inflation factors. Three sections are plotted*: *for the minimum*, *medium and maximum uplift inflation factor* $z_{\max}$ *multiplied by the standard assumed geometric profile across the section. In addition*, *the red dashed lines* (*only near the trench*, *i.e.*, *between nodes* 1 *and* 2) *display two limiting examples where the inflation factor is modified by multiplying it by the ratio* $z_t/z_{\max} \in [0, 1]$ *of either* 0 *and* 1 *in these two cases.* (c) *North–South sections of the factor* $z_{\max}$ (*fixed*) *across all series of* 14 *polygonal blocks.*

Witter et al. (2013). Application of the maximum inflation factor of 2.2 leads to seabed uplifts in our scenarios of up to 11 m, similar to the seabed uplifts of up to 10.2 m offshore from the modelled study site of Witter et al. (2013). Furthermore, similar large, localised uplifts have been identified or inferred in other subduction zones [e.g., the Nankai subduction zone; see Moore et al. (2007)] and in particular

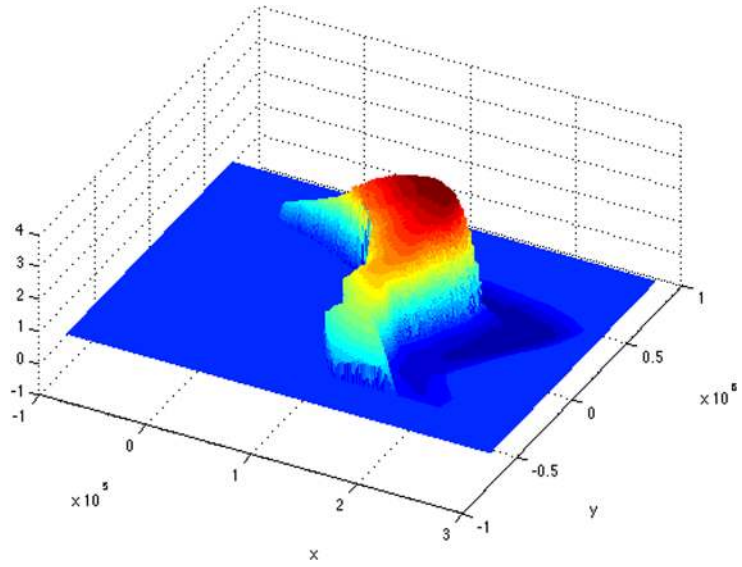

FIG. 3. *Seabed deformation* (z*-axis*) *for* $z_{\max} = 1.52$ (*i.e., maximum uplift of* 3.34 *m*), $z_t/z_{\max} = 0.67$. *All units in meters.*

giant earthquakes where they have contributed to particularly large tsunamis [e.g., Tohoku 2011; see Satake et al. (2013)].

The seabed deformation at the location of the source is added onto the initial bathymetry. The deformation shape is represented by a set of 56 polygons (into 14 rows of four arbitrary shaped 4-sided polygons; see Figure 2) designed to efficiently represent the distributions of coseismic seabed uplift and subsidence inferred to occur in Cascadia earthquakes from the incompletely known patterns of such deformation in past events such as that in 1700 AD [Atwater et al. (1995), Satake, Wang and Atwater (2003), Nelson et al. (2008), Hawkes et al. (2011), Wang et al. (2013)] as well as the "mirror image" provided by ongoing interseismic deformation [Wang et al. (1995)]. The polygons are chosen to allow the deformation to vary parallel and perpendicular to the trench (strictly, the base of the continental slope in the case of the CSZ) that marks the oceanward limit of the subduction zone fault. For the CSZ, the exact location of the landward limit of deformation is highly uncertain due to the relative lack of information about land-level movements from past events, although a variety of lines of indirect evidence can be used to constrain it [Hyndman (2013)]. Therefore, the width of the fault rupture, which is an important aspect of the rupture geometry, is highly uncertain [Nelson et al. (2008)]. Quadratic curves represent the deformation shape of the uplift and subsidence. This is a very flexible set-up with minor assumptions which avoids discontinuities at the surfaces, except as indicated by the geological and geophysical evidence. Specifically, the only discontinuity that is introduced in this

representation is at the trench, which is a realistic discontinuity where the subduction zone fault cuts the seabed [Goldfinger et al. (2012)]. Note however that the representation that we adopt allows for cases where the seismogenic fault rupture does not reach the trench (and so vertical deformation there is close to zero) as well as cases in which it does.

The middle rows of the deformation have larger uplift and subsidence heights compared to the upper and lower rows. The surfaces of neighbouring rows are smoothed to avoid discontinuities in the uplift between the polygons. This representation, using uplift and subsidence surfaces, attempts to mimic as close as possible the real source seabed deformation, and is designed to replicate evidence-based work on Cascadia [Satake, Wang and Atwater (2003), Wang et al. (2013)] while using a small number of source characteristics. The largest uncertainty in the description of the source geometry is the trench uplift and location. In order to account for this uncertainty, the trench height is allowed to vary a lot. The uplift starts from the trench location and goes to the maximum, with the location of the maximum assumed to be known at node 2; see Figure 2(b). Node 1 at the left edge is the coseismic trench location. The height at the trench is assumed to take values between zero and the maximum height (at node 2). The right edge of the uplift, located at node 3, is equal to the seabed level. The subsidence, at node 4, is assumed to be equal to 15% of the maximum uplift (located at node 2); it could vary between 0% and 30%, but 15% is considered realistic enough here.

Accurate bathymetry and topography are then combined with the seabed deformation representation to run the tsunami model VOLNA [Dutykh, Poncet and Dias (2011)] and obtain wave elevations. We run VOLNA on the GPU cluster Emerald. Figure 4 shows the dense unstructured triangular mesh comprised of 1,197,384 nodes for 2,392,352 triangles. The NOAA/NGDC bathymetry and topography data sets were converted to obtain the adaptive mesh that is used by VOLNA. Different resolutions of DEMs are merged to cover the whole CSZ ranging from 3 arc-sec to 1 arc-minute. The triangle sizes are smaller in shallow water where the waves travel more slowly, or where slopes are high, that is, where greater computational accuracy is required. This reduces computational cost while still providing accuracy where it is needed. This cost is still large, despite a huge speed-up provided by the use of multiple GPUs in the OP2 framework [Giles et al. (2011), Mudalige et al. (2012, 2013)]: around one hour of wall-clock time for two hours of simulation using 3 GPUs. To avoid reflections of simulated tsunami waves at the West and South boundaries of the domain interfering with the results produced in the time of interest for coastal inundations, a synthetic slope in the bathymetry was generated to slow down the propagation, since the speed of propagation is proportional to $\sqrt{gh}$, where g is the acceleration due to gravity and h is the water depth.

Our choice of a motion of the seabed deformation from North to South is largely arbitrary as there is no data available so far to justify or disprove it. The four polygons of each row move together, leading to a trench parallel rupture propagation as in Lay et al. (2005) and Simons et al. (2011) for the Indian Ocean 2004 and the

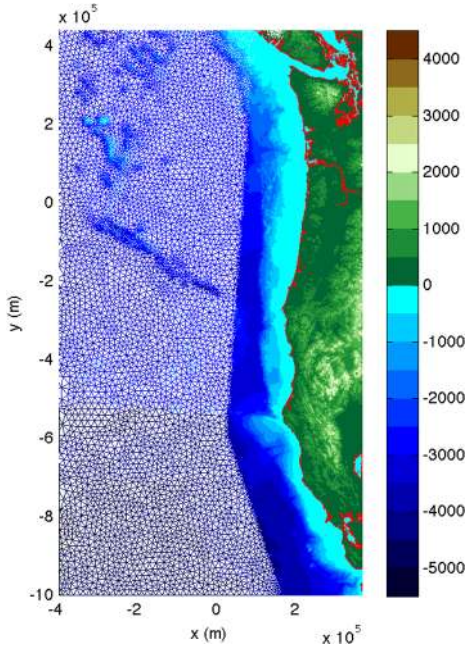

FIG. 4. *Triangular unstructured mesh of the computational domain, composed of* 2.3 *million triangles, together with elevation in meters* (*color scale*). *Note*: *the small triangles* (*down to around* 100 *m edges*) *are too dense to be visible.*

Tohoku 2011 fault ruptures respectively. The lag between the initiation of motion in the first row (North) and the last row (South) varies, with the possibility that all polygons move together simultaneously. This lag is split evenly across consecutive rows. When all the polygons complete their motion, the rupture takes its final shape. At the beginning of the motion, the polygons 1, 2, 3 and 4 in Figure 2 start to move simultaneously, followed by the next row of polygons, and so on, with the last row of polygons, numbers 53, 54, 55 and 56 the last one in motion. The shape of the seabed displacement can be observed in Figure 2 and in two movies in the Supplementary Material [Guillas et al. (2018)] for two cases described in the next section.

3. Hypothetical event cases. The primary uncertainty in tsunami predictions is induced by the imprecise description of the coseismic seabed deformation. The source is not deterministic and differs between events. For this reason, different hypothetical earthquake sources that can create tsunamis are generated by various combinations of the three source characteristics. VOLNA is run for a set of these events and the resulting tsunami wave elevation time series are obtained. The maximum source elevation factor, $z_{\max}$, takes values between 1 and 5 (it is multiplied by 2.2, corresponding to a physical range of 2.2 m to 11 m), and the ratio $z_t/z_{\max}$ between 0 and 1, as justified by the discussion and references in the previous section

and in Wells and Coppersmith (1994). The ratio of the whole rupture propagation duration t_p from North to South, to the reference duration $t_r = 150$ s, t_p/t_r, varies from 0 to 6, that is, from instantaneous rupture propagation to 900 s. Such range reflects the rupture propagation times for the Indian Ocean 2004 (> 1000 s, but mostly in the first 500 s), Chile 2010 (~ 140 s) and Tohoku 2011 (~ 200 s) events; see Lay et al. (2005), Lay et al. (2010) and Simons et al. (2011) respectively for these three events.

The 40 combinations of the three input characteristics cover the 3D input space using a maximin Latin Hypercube Design (LHD). Note that other techniques could be used to generate more efficiently a set of computer experiments, see Beck and Guillas (2016) for new sequential designs, with an illustration to tsunami modelling; there is little benefit here as we do not need many runs and these are relatively affordable. Each combination corresponds to a different hypothetical seabed deformation event. VOLNA simulations produce tsunami wave elevation time series at each of the 2.4 million triangles in the computational domain, a massive output. We have looked at many spatial locations to examine the outputs, but decided to analyze the wave elevation time series at two gauges only, shown in Figure 5(a): one near the entrance of the Strait Juan de Fuca, offshore of Vancouver Island, BC, Canada (gauge 141) and one near shore the city of Victoria, BC, Canada (gauge 32,298). Indeed, these two gauges are in areas that may be affected by a tsunami, and are emblematic of the types of wave forms that are created over the region.

A smooth, quasi-periodic train of waves is created near the source (gauge 141). The same tsunami wave train enters the strait and propagates over the complex bathymetry. It also bounces off the boundaries of the strait and travels around islands. The resulting complex wave (gauge 32,298) differs strongly from the original wave since it displays the signature of this propagation: it starts with a trough, as observed in many past tsunami events very near shore, displays some kinks and a sharp rise, with a second wave much reduced compared to the first; see Figure 5. These two locations provide an excellent testbed for our approach: the gauge 32,298 will showcase the benefits of our approach as it requires complex data-driven basis functions, with time alignment of the runs, as explained in the next sections; the gauge 141 only confirms that there is no loss in using our complex method on a simple wave form.

VOLNA uses as inputs the computational mesh, the bathymetry and topography, and the various dynamic seabed deformations. The comparisons between the different combinations of the three source characteristics at gauges 141 and 32,298 are presented in Figure 5. Obviously, the wave amplitude increases significantly with $z_{\max}$. For a similar $z_{\max}$, the larger $z_t/z_{\max}$, the larger is the wave elevation. This indicates that when there is a larger discontinuity at the trench, higher wave amplitudes are obtained. The duration of source propagation does not seem to influence crucially the tsunami wave amplitude whereas it significantly affects the wave propagation time, as it moves the location of the maximum elevation (peak) of the time series. The smaller t_p/t_r, the faster the wave arrives at the gauge.

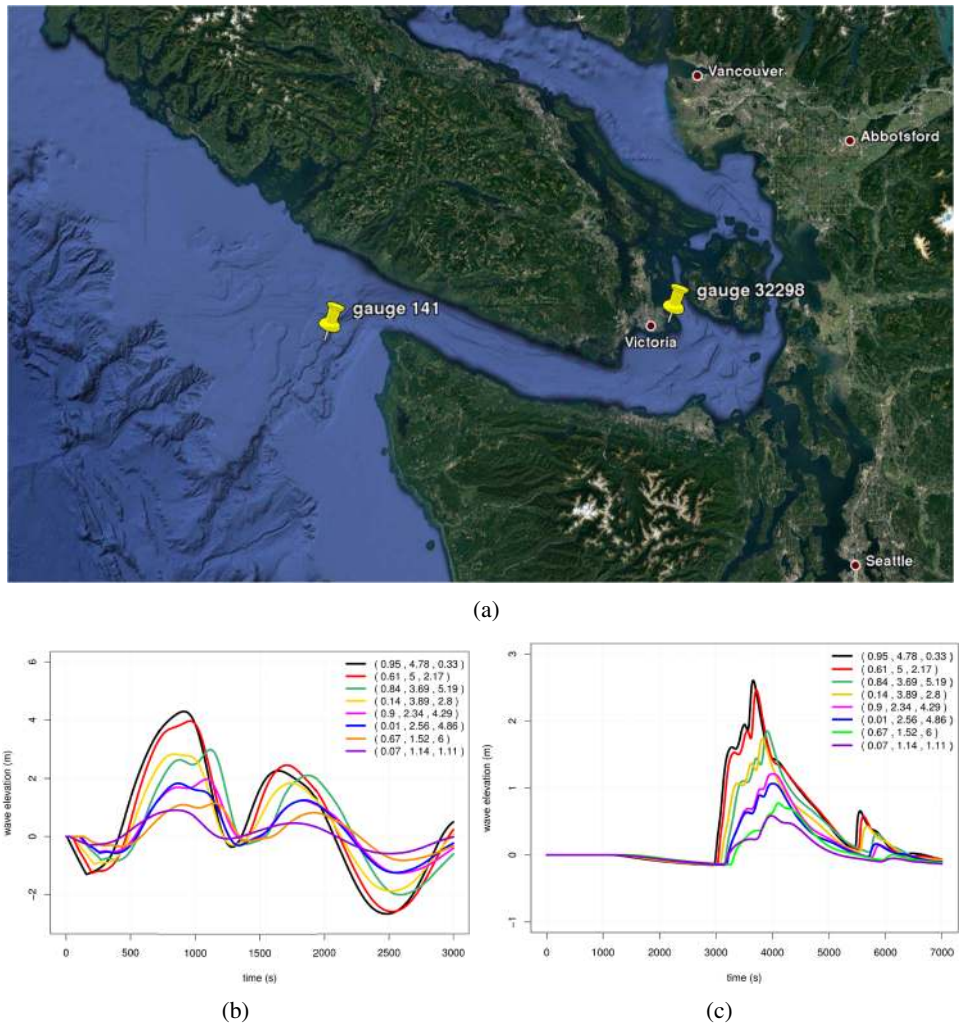

FIG. 5. (a) *Gauges locations*. *Simulated wave elevations at two gauges* 141 [*panel* (b)] *and* 32,298 [*panel* (c)] *with corresponding color and values of the seafloor deformation parameters* ($z_t/z_{\max}$, $z_{\max}$, t_p/t_r) *in the legends*.

4. Functional representation of the outputs. We introduce here a novel functional representation for the outputs of computer models. Such approaches have been pursued before with a fixed basis [Bayarri et al. (2007), Hung, Joseph and Melkote (2015)]. We specifically introduce the data-driven technique of Functional Principal Components, combined with the functional registration as a preprocessing step. To our knowledge, this is the first time that FPCA and registration are used in emulation. Functional data are data objects consisting of functions varying over a continuum, which is usually time. They are observed as n discrete

pairs (t_j, y_j), $j = 1, \ldots, n$, with y_j the function value at time t_j. It is assumed that there exists a smooth (randomly drawn) function f that can produce the observed data. In that situation, continuity of f implies that pairs of adjacent outputs, y_j and y_{j+1} are similar to each other; smoothness entails further that these local variations are small, as typically first and second derivatives are assumed to be themselves varying slowly.

Our functional modelling proceeds as follows: (1) we apply functional registration to all curves to eliminate phase variations across curves; (2) we select traditional cubic splines as basis functions for the representation of each individual time series of output; (3) FPCA [Ramsay and Silverman (2005), Ramsay, Hooker and Graves (2009)] is employed to obtain the main modes of variation that describe the time series. FPCA has been shown to provide more meaningful modes of variability than multivariate PCA [Ramsay and Silverman (2005)] as the underlying physics smoothly drives the variability, whereas in large dimensions PCA can yield numerical instabilities in the analysis for instance due to collinearities and small gaps between eigenvalues [Castro, Lawton, and Sylvestre (1986), Jolliffe (2002)].

Certain features of the simulated waves, such as peaks, do not occur at the same time for different durations of the source propagation. As a result, the variations in amplitude of the wave resulting from changes in the other inputs are obfuscated. Thus our ability to propagate uncertainties or examine sensitivity of the wave to these inputs is damaged; we need to rectify this problem by aligning the times at which curves vary, a technique also called registration. Note that the maximum wave elevation is a very important measure of impact on structures and people. Without registration we would clearly underestimate the potential for high values.

Landmark registration [Ramsay and Silverman (2005)] is chosen here since visible landmarks exist across outputs and are clear. The transformations, using warping functions, capture and eliminate phase variations of the original data set. To do so, the first step is the selection of landmarks used to align the curves, by estimating strictly increasing transformations of time, $h_i(t)$ for each i in the LHD. For a time interval $[0, T]$, time warping functions must satisfy the general constraints $h(0) = 0$ and $h(T) = T$. The registered functions are then described with the equation $f_i^*(t) = f_i[h_i^{-1}(t)]$. Note that alternative forms of registration exist and can be even combined with FPCA [Kneip and Ramsay (2008)]. Landmark registration is rather arbitrary, and cannot be automated to a set of gauges as these are clearly gauge-dependent. The problem of finding automated forms of registration, especially when features, such as peaks or crossings, do not repeat across runs (e.g., when a wave sometimes reaches a gauge, sometimes not), is beyond the scope of this paper.

The results after applying landmark registration are satisfactory. Indeed, as seen later in Leave-One-Out diagnostics for the whole procedure, this choice of registration leads to an almost perfect match—especially in the finest details of peaks and kinks—between simulated and emulated outputs. For gauge 141, two landmarks are selected: the locations of the second and third minima; see Figure 5(b). For

gauge 32,298, two landmarks are also selected: the locations of the first maximum and first minimum of the time series; see Figure 5(c). The landmark registration will provide additional precision in the tsunami amplitudes (especially for gauge 32,298) where variations in wave elevation are sharp.

FPCA is then applied to the registered functional data. The resulting modes of variation, known as FPCs or harmonics, will then be used as basis functions g_j^s for the outputs in (1): $g_j(r,s) = g_j^r(r) \otimes g_j^s(s)$. The advantage of FPCA over multivariate PCA is that it takes into account the smoothness of the data. Data-driven (empirical) basis functions are then defined as the FPCs and output functions are expressed according to their coefficients (or scores) in this basis. In more detail, using a roughness penalty approach, we first estimate the coefficients of expansion of the time series $f_i^*(t)$ into a B-spline basis ϕ_k, $k = 1, \ldots, K$. We obtain a function that approximates f_i^*, $\xi_i(t) = \sum_{k=1}^K c_{ik}\phi_k(t)$, by minimizing the penalized least squares criterion with a smoothing parameter λ. The smoothing parameter λ, chosen empirically by generalised cross-validation (GCV) in (4), controls the level of smoothing since the second derivative measures the roughness of the fitted curve:

$$\sum_{j=1}^{n}\left[f_i^*(t_j) - \sum_{k=1}^{K} c_{ik}\phi_k(t_j)\right]^2 + \lambda \int \left[\xi_i''(t)\right]^2 dt. \tag{4}$$

FPCs are obtained by incorporating a penalty term in the orthonormality constraint imposed on the traditional PCs, and hence are smooth. We used the same B-spline basis to represent the PCs as $\zeta_i(t) = \sum_{k=1}^K z_{ik}\phi_k(t)$, and also measured roughness by the integrated squared second derivative of the function. After projecting the PCs onto a B-spline basis, the fitting criterion for the coefficient vectors $\mathbf{z}_i$ of the FPCs becomes [Silverman (1996)]:

$$\max \frac{\mathbf{z}_i^T \mathbf{\Phi}^T \hat{\mathbf{Y}}^T \hat{\mathbf{Y}} \mathbf{\Phi} \mathbf{z}_i}{\mathbf{z}_i^T (\mathbf{I}_K + \lambda \mathcal{P}^T \mathcal{P}) \mathbf{z}_i}, \tag{5}$$

where $\hat{\mathbf{Y}} \in \mathbb{R}^{K \times n}$ denotes the matrix consisting of the centered fitted values of wave elevation from the smoothing splines, and $\mathbf{\Phi} \in \mathbb{R}^{n \times K}$ denotes the B-splines basis matrix composed of (j,k) entries $\phi_k(t_j)$. $\mathcal{P}^T\mathcal{P}$ is the $K \times K$ penalty matrix with entries $\int \phi''_m(t)\phi''_l(t)\,dt$, and thus $\mathbf{z}_i^T \mathcal{P}^T \mathcal{P} \mathbf{z}_i$ quantifies the roughness of each PC. We applied the same smoothing parameter λ to obtain all PCs. Note that the numerator is the variance between $\hat{\mathbf{Y}}$ and ζ_i to be maximized, as PCs have that goal, and the denominator is the orthonormal constraint imposed on the PCs, albeit with added smoothing penalty here. Once $\hat{\mathbf{z}}_i$ are estimated from (5) we obtain the PC by $\hat{\zeta}_i(t) = \sum_{k=1}^K \hat{z}_{ik}\phi_k(t)$. The R-function `pca.fd` in `FDA` library is used to carry out the estimation.

For both gauges, the first two PCs are enough to describe the variations; see Figures 6, 7 for registered data. Since the principal components represent variations around the mean, the mean curve is plotted along with the $+$'s and $-$'s

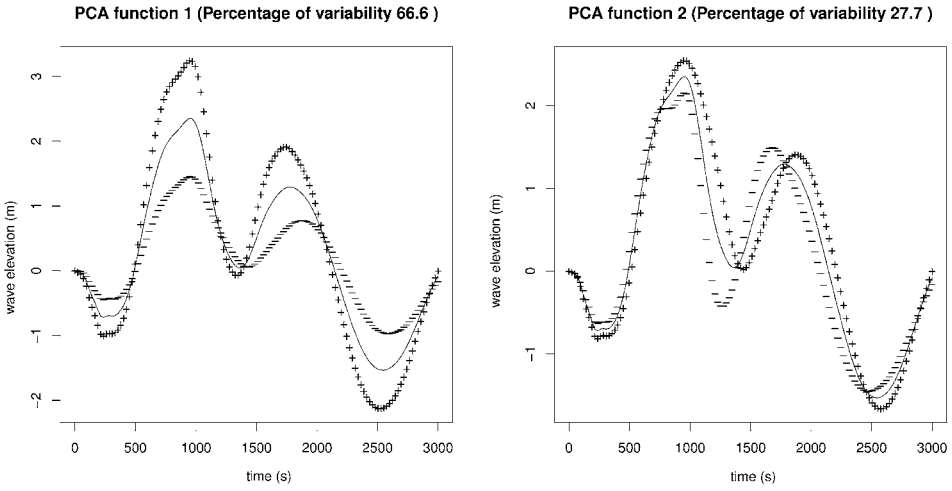

FIG. 6. *First two PCs, registered gauge* 141. *Solid line*: *mean wave elevation*, ± *lines*: *variability along the PC directions. These indicate how much and in which direction the wave form varies when the PC loading is increased.*

that correspond to the consequences of adding and subtracting an amount of one standard error in the direction of each principal component. Note that if FPCA is carried out on the unregistered data, the first principal components account for a lower percentage of variation, immediately demonstrating the benefit of using registration. For gauge 141, the first PC displays large variations at the maxima and

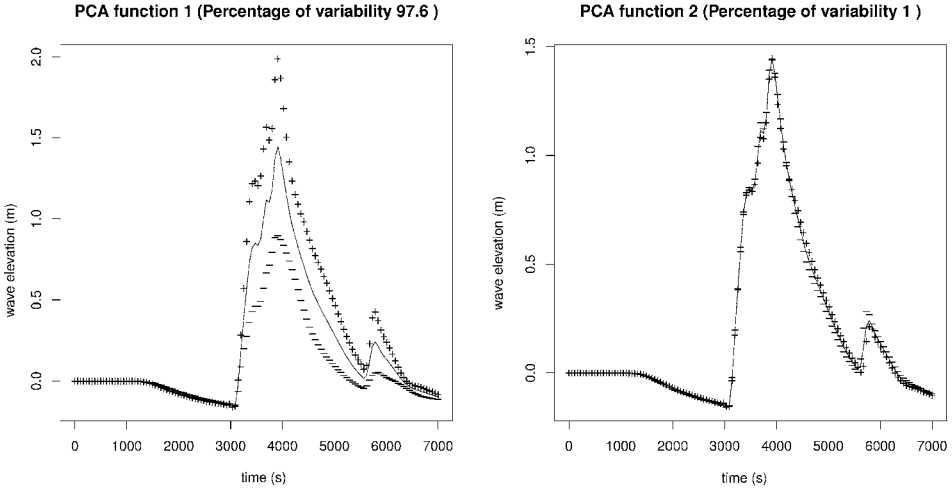

FIG. 7. *First two PCs, registered gauge* 32,298. *Solid line*: *mean wave elevation*, ± *lines*: *variability along the PC directions. These indicate how much and in which direction the wave form varies when the PC loading is increased.*

minima. Despite the registration, the second PC for gauge 141 still shows phase variations after 1500 s (note that the first PC does not show phase variations). It is likely that registration did not perform well in this time window for the second PC, possibly because of the large amount, and complexity, of phase variations. It is also likely the the variations in amplitude are mostly taken into account by the first PC so phase variations left over are greater than any other amplitude variations: here the signal is rather difficult to understand physically. The first two PCs capture 94.3% of the variations; the second PC accounts for a significant percentage of the variations (27.7%). For gauge 32,298, the first PC accounts for 97.6% of the variability, and shows a very small variability initially up to time 3000 s. This is expected as the first incoming wave arrives after 3000 s. There are large variations near the peaks. The second PC represents only 1% of the variation and is not informative, unlike for gauge 141, as almost all the variability is captured by the first PC, since the wave elevation variations across runs is not influenced anymore by the ratio at the trench $(z_t/z_{\max})$ after traveling some distance from the source.

5. Statistical emulation of the tsunami model. The uncertainties in tsunami sources make Probabilistic Tsunami Hazard Analysis (PTHA) necessary [González et al. (2009)]. However, PTHA requires ideally a large number of high-resolution simulations (e.g., using VOLNA). To make this feasible, two typical approaches are pursued. One uses a large number of runs of a simplistic model, often at low resolution. The other one uses only a small number of high-resolution runs. A statistical emulator is built to overcome the computational problem of generating a large number of runs and obtain fast predictions used for the final step of PTHA. Specifically, an Outer Product Emulator is used here with four choices of basis functions for the outputs of tsunami wave elevation time series: Fourier [Sarri, Guillas and Dias (2012)] or FPCs are used for the output regression functions, both with or without registration. It illustrates the benefit of using registered FPCs as proposed in the previous section.

5.1. *Building an emulator.* Model evaluations are available for 40 combinations of the three source characteristics $r = (z_t/z_{\max}, z_{\max}, t_p/t_r)$, with each one representing a hypothetical fault rupture that generates a tsunami. These 40 points are chosen to cover the three-dimensional input space using a maximin LHD, within the following domain: $\frac{z_t}{z_{\max}} \in [0,1]$, $z_{\max} \in [1,5]$ and $\frac{t_p}{t_r} \in [0,6]$ (we normalize the input intervals over $[0,1]$).

For the emulation, appropriate choices of regression and covariance functions for inputs r and output argument s, where s is time, are made. The set of input regression functions, $G^r \stackrel{\text{def}}{=} \{g_1^r, \ldots, g_{\nu_r}^r\}$, where ν_r is the number of input regressors, contains polynomials for each of the three input parameters. We employ seven input regressors: constant, linear and quadratic polynomials. These regressors are orthogonal Legendre polynomials, shifted into the unit integral [0, 1];

see (22) in Sarri, Guillas and Dias (2012). The main modes of variability in the input-output relationship are mostly captured by the regression function and the covariance structure is flexible enough to model the rest, in case the nonlinearity is more complex than quadratic. Note that these modes turned out to be simple: the influence of parameters, as seen eventually in the sensitivity analysis below (for the maximum elevation only though), is actually mostly linear. The regressors in inputs are always the same for all types of output representation (Fourier or FPCs). Tensor products of Legendre polynomials for emulation have also been used successfully in another emulation context [Kaufman et al. (2011)] using compactly supported correlation functions (reducing dramatically computing costs) to model small-scale variability, and regression functions in the mean of the GP to model large-scale variability (up to fifth order Legendre polynomials as a result of a restricted family of correlation functions). As we did not restrict ourselves to compactly supported correlations functions, our choice of second order Legendre polynomials constitute a reasonable trade-off in our setting.

For the output regression functions, we define $G^s \stackrel{\text{def}}{=} \{g_1^s, \ldots, g_{\nu_s}^s\}$, where ν_s is the number of output regressors. We consider two choices of regressors: (1) Fourier terms in addition to a constant term; (2) functional PCs as explained above. We now justify the choice of Fourier terms. On the one hand, we want to represent the time series with Fourier terms whose frequencies reflect the possible time series of outputs. On the other hand, having too many frequencies leads to an increase in computational costs and noisy variations, especially with high frequencies. Hence the frequencies are different for each gauge as they oscillate differently. We inspected the frequencies of each wave time series for the different inputs and fitted these outputs using several sets of Fourier modes. After trial-and-error investigations, a set of frequencies that leads to good emulated predictions is selected for each gauge. Specifically, for gauge 141 the set of frequencies ω is

$$\Omega = \left\{ \frac{1}{1600}, \frac{1}{1400}, \frac{1}{1200}, \frac{1}{1000}, \frac{1}{800}, \frac{1}{600}, \frac{1}{400} \right\}.$$

Hence, the set of output regression functions is

$$G^s = \{1, \sin(2\pi T/\omega), \cos(2\pi T/\omega)\}, \qquad \omega \in \Omega.$$

The set of frequencies for gauge 32,298 is

$$\Omega = \left\{ \frac{1}{6400}, \frac{1}{5800}, \frac{1}{5200}, \frac{1}{4600}, \frac{1}{4000}, \frac{1}{3400}, \frac{1}{2800}, \frac{1}{2200}, \frac{1}{1600}, \frac{1}{1000}, \frac{1}{400} \right\}.$$

Indeed, as seen in Figure 5, the time series for gauge 32,298 clearly displays high frequency variations that require the use of a different set of frequencies from gauge 141. The frequencies are the same for registered and nonregistered data as registration changes infinitesimally the frequencies of oscillations in the curves. The number of selected Fourier basis functions can become large in order to capture various types of oscillations. FPCs as time input regression functions should

give better results due to their ability to directly pick the right modes of variation across curves.

The input and output residuals covariance functions, κ^r and κ^s, are chosen to be the widely used power exponential functions:

$$\kappa^r = \exp\left(-\left(\frac{|\frac{z_t}{z_{\max}} - \frac{z_t}{z_{\max}}'|}{\lambda_1}\right)^{3/2}\right) \times \exp\left(-\left(\frac{|z_{\max} - z'_{\max}|}{\lambda_2}\right)^{3/2}\right) \times \exp\left(-\left(\frac{|\frac{t_p}{t_r} - \frac{t_p}{t_r}'|}{\lambda_3}\right)^{3/2}\right), \tag{6}$$

$$\kappa^s = \exp\left(-\left(\frac{|t - t'|}{\lambda_t}\right)^{3/2}\right), \tag{7}$$

where λ_1, λ_2, λ_3 are the correlation lengths for the inputs and λ_t for the outputs. The $3/2$ power constitutes a good trade-off between no differentiability of a GP equipped with a correlation above with a power equal to 1 (though still not differentiable with power equal to $3/2$ but not as rough in practice) and infinite differentiability for a value of 2. It was also selected by Kaufman et al. (2011) after careful investigation of these choices of power and their effects but in another context (practitioners often use 1.9 or 2). The influence of the smoothing parameter is relatively small between 1.5 and 2 and cannot explain the predicted interval widths seen in Figures 8 and 9. These widths, for registered functional data (a necessary step for these data) are fine, as even a slightly reduced width would result in much more than 5% of the data outside these intervals (when recasting the problem pointwise empirically, not exactly the case here but an informative view nevertheless). For gauge 141, when the best modelling steps are followed (registration and PCs), the empirical pointwise coverage is probably slightly above 95%, but would reduce by a large amount with even a slightly narrower width. Clearly predictions at the peak would lie outside these intervals for gauge 32,298.

The correlation lengths, after initial exploration of a reasonable range, were selected using a simple trial and error method, using different combinations of $\lambda_1 = \lambda_2 = \lambda_3 = 0.1, 0.5, 0.9, 1, 2$ (as little gain is obtained in using different ones) and $\lambda_t = 0.1, 0.5, 0.9$ for all four cases (Fourier or FPCs, unregistered or registered) and all gauges. The resulting emulator's predictions are validated and compared using the total mean CI length and RMSE. We estimate a set of correlation lengths and use these values everywhere for all techniques. The selection of the correlation lengths could be performed in a more refined manner using marginal likelihood, as in our previous work [Sarri, Guillas and Dias (2012)]. However, this entails more computations, for little additional benefit. Indeed, the subsequent predictions below using these correlation lengths are very good. Besides, the dependence to correlation length is weak and thus will not yield a significant improvement; see Rougier et al. (2009) for a sensitivity analysis of the OPE to correlation lengths. The final selection is $\lambda_1 = \lambda_2 = \lambda_3 = 1$ and $\lambda_t = 0.5$.

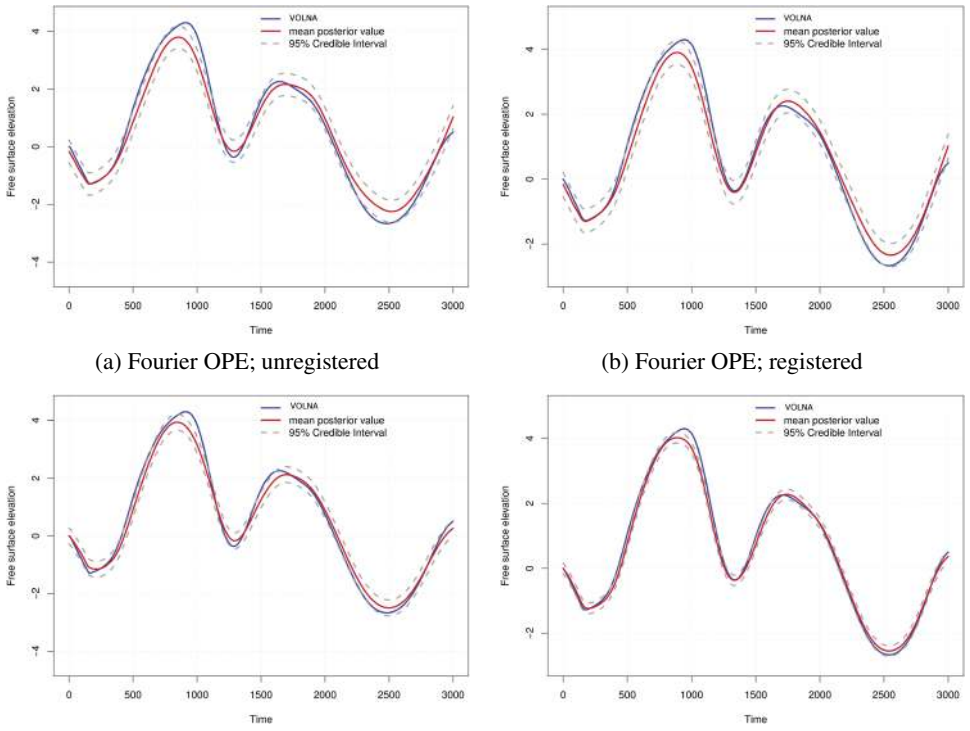

(a) Fourier OPE; unregistered

(b) Fourier OPE; registered

(c) FPCs as OPE; unregistered

(d) FPCs as OPE; registered

FIG. 8. *LOO diagnostics plots, gauge* 141. *Units: elevation in meters and time in seconds. Source* $z_{\max} = 4.78$ *m,* $z_t/z_{\max} = 0.95$ *and* $t_p/t_r = 0.33$ (*LHD point* 5). *Simulated wave in blue, mean emulated wave in red*; 95% *credible intervals display the* 95% *most likely emulated wave forms.*

The last step is the choice of hyperparameters $\{m, V, a, d\}$, done here as in Sarri, Guillas and Dias (2012). With this selection of parameters, the prior emulator is constructed and combined with the model evaluations. The resulting OPEs are able to give quick probabilistic predictions of the wave elevation for any combination of the three source characteristics. A Leave-One-Out (LOO) validation is used to assess accuracy of the emulation. 40 LOO diagnostic plots are obtained for each gauge in the four cases. For most of the cases, the LOO diagnostics show that the emulator predictions are satisfactory, with both the shape and the peaks accurately predicted. The 95% credible intervals capture well the simulations. We noticed that the set of 40 curves used in LOO lie almost entirely within these bands, showing that the 95% coverage probability is seemingly correct in all cases, but that the the CIs might be all too generous. The mean credible interval length (MCIL) and the root-mean-square error (RMSE) between the observed and the predicted values for each of the 40 inputs are calculated and used to assess the predictions.

LOO diagnostic plots at the same LHD point (05) are displayed in Figure 8 for gauge 141 and in Figure 9 for gauge 32,298. For gauge 141, the use of FPCs in

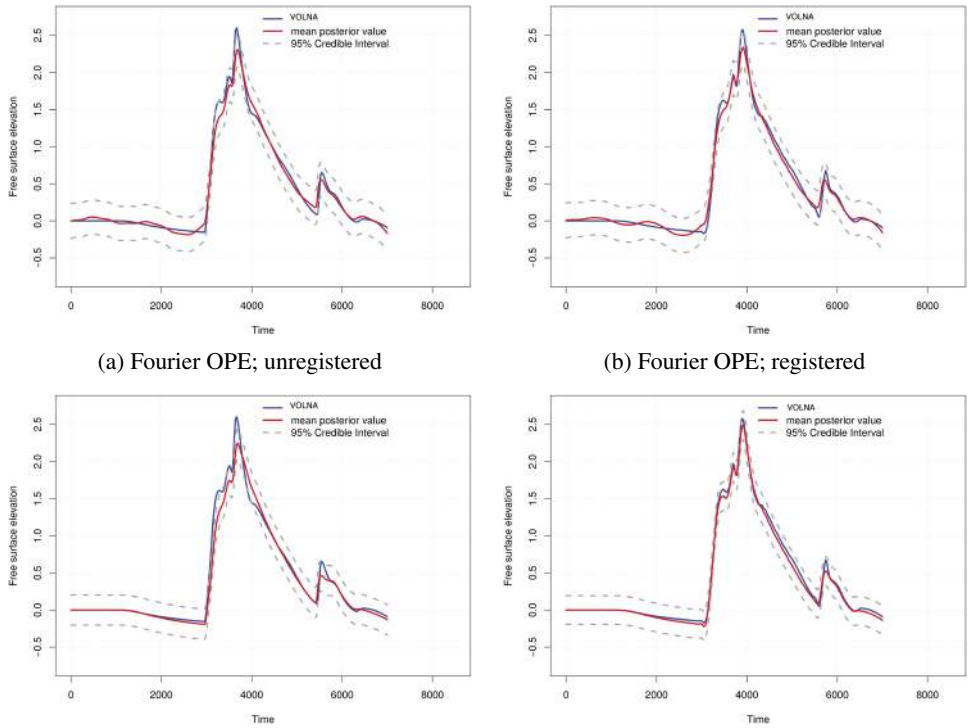

(a) Fourier OPE; unregistered

(b) Fourier OPE; registered

(c) PCs as OPE; unregistered

(d) PCs as OPE; registered

FIG. 9. *LOO diagnostics plots, gauge* 32,298. *Units*: *elevation in meters and time in seconds*. *Source* $z_{\max} = 4.78$ *m*, $z_t/z_{\max} = 0.95$ *and* $t_p/t_r = 0.33$ (*LHD point* 5). *Simulated wave in blue, mean emulated wave in red*; 95% *credible intervals display the* 95% *most likely emulated wave forms*.

place of Fourier basis functions improves the predictions for both unregistered and registered cases, since the wave elevation oscillations at gauge 141 are smooth, without sharp fluctuations. The complex shape makes the emulator's job relatively hard for gauge 32,298. This gives the opportunity to observe significant differences between the predictions of the four emulations. The specific LHD point has a ratio of the trench height to the middle maximum height, $z_t/z_{\max}$, close to 1. This indicates a large discontinuity at the trench. Also, $z_{\max}$ is close to the upper limit of the range. Hence, the resulting tsunami waves have large amplitudes and variations, which makes the emulator's job even more difficult. Using Fourier basis functions, the emulator predictions cannot escape from oscillating even at locations where the wave elevation is almost flat, which is the case for the initial 3000 s. However, the emulator performs rather well even with the unregistered data and the selection of Fourier basis functions. The Fourier approach yields variations at the frequencies required, as seen in the similar kinks occurring at around 3500 s in the registered case (one is actually captured, the other not, but both are in the FPC case). From

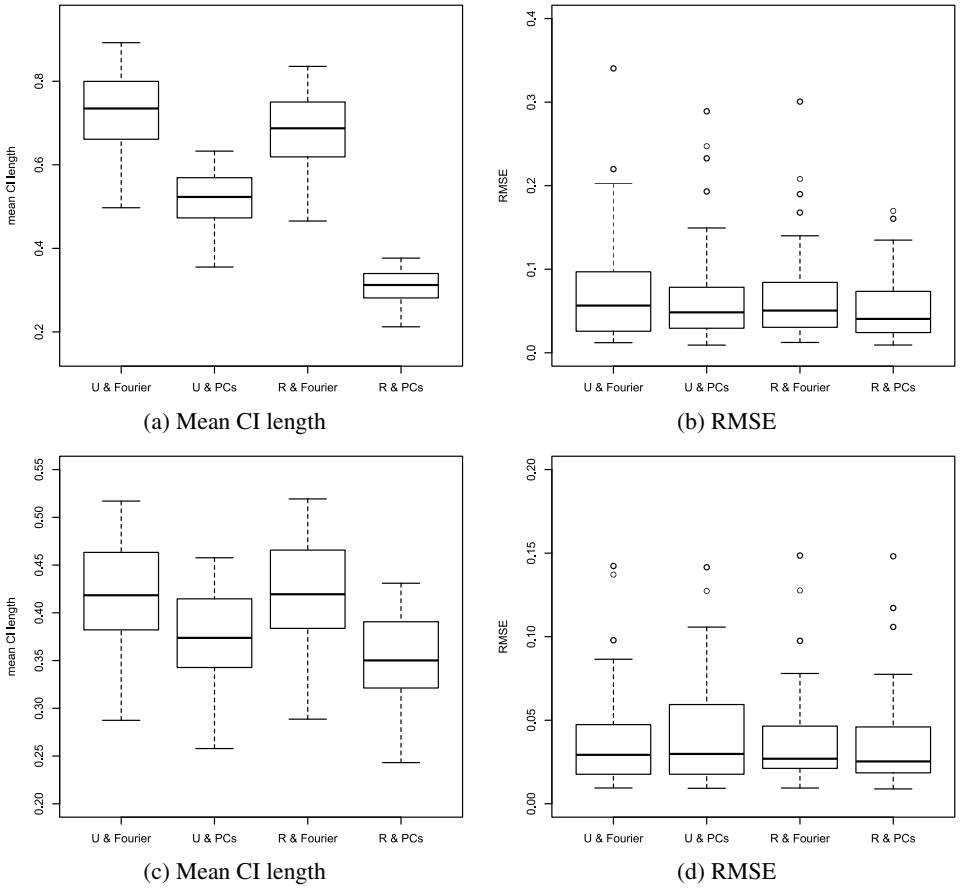

(a) Mean CI length (b) RMSE

(c) Mean CI length (d) RMSE

FIG. 10. *Total Mean Credible Interval length and RMSEs for OPEs*, *gauges* 141 (*upper panels*) *&* 32,298 (*lower panels*). *R*: *registered*; *U*: *unregistered*.

Figure 9 it is obvious that the use of FPCs does not display spurious oscillations, as opposed to Fourier terms. Note that in the case of the unregistered data, the use of FPCs is unable to catch the details in the shape as well as for registered FPCs.

Figure 10 displays the comparisons of pointwise mean CI lengths and RMSE statistics across the four cases, computed using the 40 LOO diagnostics. What is a coverage probability for random curves is not yet possible to answer formally or even numerically. Indeed, the notion of a curve being (not necessarily entirely) within a credible interval is not clear (as opposed to a pointwise case), but may be in the future using either joint prediction intervals or Functional Data Analysis notions of data depth for instance [López-Pintado and Romo (2009)]. Overall, the use of landmark registration together with FPCs mainly improves the CI lengths, with limited improvements in the RMSEs. Indeed, the meaningful errors, as seen in LOO diagnostic plots, are at the peaks, but these reductions are tiny in terms of

RMSEs as these are computed over the whole time series; see, for instance, panels (c) and (d) in Figure 9 where the improvement in terms of the peak estimation (over the small interval 3800–4000 s) by the registration is offset by a lack of fit over the longer period 4500–5500 s. In conclusion, the use of landmark registration together with FPCs is the best method overall. Importantly, the FPC-based OPE captures better the peaks of the waves, a critical value to evaluate potential impact of tsunamis.

5.2. *Sensitivity analysis*. A sensitivity analysis would be extremely expensive using a Monte Carlo approach without the help of an emulator. Measures of sensitivity have been developed in the past. Oakley and O'Hagan (2004) for instance provide graphical representations of posterior expectations of the individual conditional expectations of the response with respect to each input. Our graphical outputs are very similar, but instead of computing the posterior expectation to capture an overall effect, we break down the response by the value of the other two coefficients, as we can afford this computation in our 3D case.

Figure 11 shows the sensitivity of gauge 141 outputs to individual inputs. For $z_{\max}$, emulator predictions are obtained at combinations of 0.01 increments of the parameter $z_{\max}$ ($z_{\max} = [1, 1.01, 1.02, \ldots, 5]$) and for the maximum, middle and minimum values of the range of the two other parameters, $z_t/z_{\max}$ and t_p/t_r. This amounts to 3609 different combinations. Similarly, we run 909 combinations of 0.1 steps for $z_t/z_{\max}$ and the maximum/middle/minimum values of $z_{\max}$ and t_p/t_r. We also run 5409 combinations to examine the sensitivity to t_p/t_r for some choices of the two other parameters ($t_p/t_r = [0, 0.01, 0.02, \ldots, 6]$, $z_{\max} = [1, 3, 5]$ and $z_t/z_{\max} = [0, 0.5, 1]$). Figure 11 identifies, as expected, a strong positive correlation between the maximum wave elevation and $z_{\max}$. Also, there is a slight increase in the wave elevation, when the trench height z_t increases closer $z_{\max}$. No significant dependency on the rupture propagation time is observed.[3] Similar dependencies are valid for the gauge 32,298 (not shown). In a sensitivity analysis to combinations of two inputs using the emulator (not shown), we found a small two-way interaction. The lack of sensitivity and low sensitivity, respectively to time and trench height, cannot be assumed to be true for all possible gauges (e.g., a gauge over the trench would show a clearer sensitivity to trench height).

5.3. *Uncertainty analysis*. The uncertainty analysis propagates the uncertainty in the input characteristics $z_{\max}$, $z_t/z_{\max}$ and t_p/t_r to the maximum wave elevation. Using expert judgement about the most probable values for these inputs, a distribution is assigned to each of these: 2000 random input combinations are drawn, and the emulator readily provides maximum wave elevations for these

[3]The variation in rupture propagation time that we used is small compared to the tsunami wave period (i.e., rupture propagation is fast compared to tsunami propagation, as observed in fault rupture-sourced events but landslides may be a different matter).

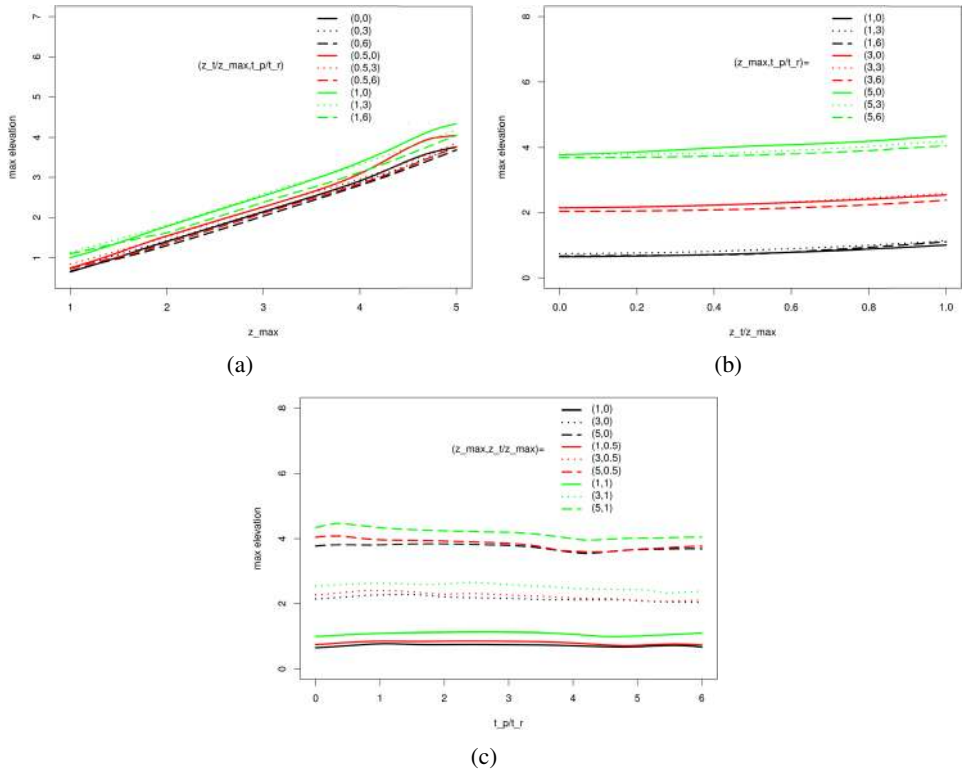

FIG. 11. *Sensitivity of the maximum wave elevation* (*in meters*), *gauge* 141, *to* (a) $z_{\max}$ *inflation factor*, (b) $z_t/z_{\max}$ *and* (c) t_p/t_r.

2000 cases. The first parameter $z_{\max}$ is assumed to have a Beta distribution Be(2, 3) over [2, 5], not over [1, 5] as we expect a maximum seabed deformation of at least $2.2 \times 2 = 4.4$ m when a fault rupture occurs in this hazard assessment. Indeed, the parameter $z_{\max}$ relates closely to the seismic magnitude, and to the coastal subsidence as Goldfinger et al. (2012) argue that the coseismic uplift relates to the proxy of turbidite volume. Leonard et al. (2010) describe the distribution of the coastal subsidence and hence help define the distribution of $z_{\max}$. The dependency of $z_{\max}$ to the coastal subsidence is particularly valid for cases where the ratio of the maximum uplift to the maximum subsidence is considered constant, which is the case here. For the whole margin rupture case, the values of $z_{\max}$ are considered to be more likely between 2 and 4, and less likely between 4 and 5; see Figure 12. This range of values for $z_{\max}$ is supported by Wells and Coppersmith (1994), who observe that there is a range of slips associated with a fixed fault rupture length that does not go to zero.

We assume that $z_t/z_{\max}$ follows a truncated Normal distribution, with underlying Normal mean 1 and variance 0.09, truncated over [0, 1]; see Figure 12. Indeed, $z_t/z_{\max}$ is assumed to be in the range [0, 1]. Note that a ratio larger than 1 would be

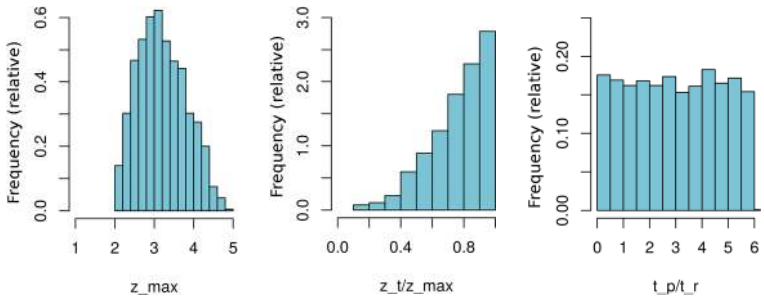

FIG. 12. 2000 *draws from the distributions of* (*dimensionless*) *parameters* $z_{\max}$, $z_t/z_{\max}$, t_p/t_r.

conceivable. However, since for our analysis the whole rupture propagation case is investigated, we are not considering that this ratio can take values larger than 1 all along the margin, and thus we restrict ourselves to [0, 1]. There is support for the fact that $z_t/z_{\max}$ should be close to 1 [Masterlark and Hughes (2008)], which is assumed here. For the last input, the limited knowledge about the duration (and even direction) of propagation leads to the selection of Uniform distribution between 0 and 6. 2000 random samples are drawn from the three distributions; see Figure 12. Note that physically, if the uplift size $z_{\max}$ is large, then the discontinuity at the trench $z_t/z_{\max}$ is likely to be large [Wells and Coppersmith (1994)]. We create a sample of non-Gaussian variables that respects this fact, by sorting the values of these two parameters alongside and by merely associating large values of one to large values of the other. These combinations are combined with the values of t_p/t_r in random order to generate 2000 different coseismic source deformation events. Emulation is employed for these 2000 different input combinations. For each gauge, the maximum elevation and the mean predictive Credible Interval (CI) length have been estimated for each combination of the input parameters. The variations are quantified using quantiles; see Table 1. The distributions of the resulting maximum wave elevations are shown in Figure 13. This analysis can be repeated for any location to assess risk over entire regions. The total time

TABLE 1
Maximum predicted wave elevation and mean predictive Credible Intervals (*CI*) *length percentiles for gauge* 141 (*first two rows*) *and gauge* 32,298 (*last two rows*), *resulting from the uncertainties drawn in Figure* 12

	0.1%	**1%**	**5%**	**25%**	**50%**	**75%**	**95%**	**99%**	**99.9%**
Max elevation	1.42	1.48	1.66	2.12	2.52	2.93	3.50	3.82	4.14
Mean CI length	0.22	0.31	0.39	0.55	0.64	0.71	0.77	0.79	0.80
Max Elevation	0.82	0.92	1.07	1.35	1.59	1.80	2.02	2.15	2.42
Mean CI length	0.10	0.13	0.17	0.24	0.27	0.31	0.33	0.34	0.34

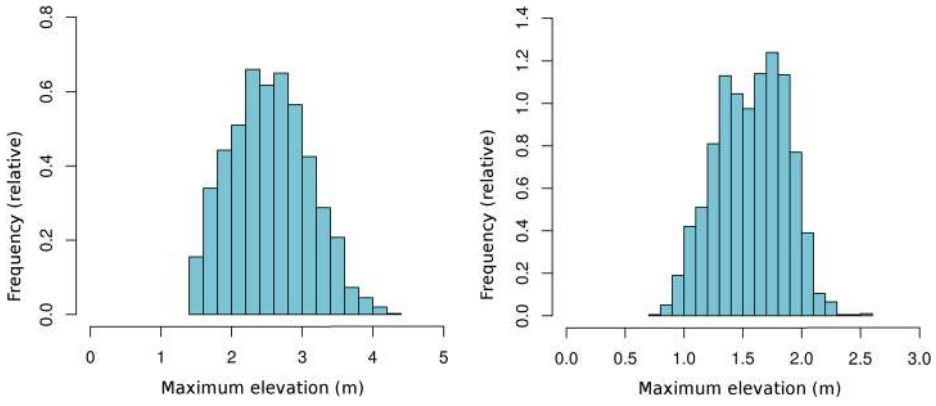

(a) Gauge 141 (b) Gauge 32,298

FIG. 13. *Maximum wave elevation predictions distribution for* (a) *gauge* 141 *and* (b) *gauge* 32,298, *resulting from the uncertainties drawn in Figure* 12.

required was of the order one minute (on a standard desktop) for 2000 emulator predictions. This highlights one more time the importance of replacing expensive computer models with accurate statistical surrogate models.

6. Conclusion. Our functional emulator efficiently and parsimoniously approximates tsunami wave forms, making use of landmark registration and Functional Principal Component Analysis. It outperforms previously used functional emulators employing a fixed basis. The emulation allows the investigation of sensitivity to input parameters describing realistic features, as well as the propagation of uncertainties for high resolution tsunami risk assessments. As we produce entire time series of uncertain tsunami waves, other summaries than the maximum elevation presented here could be obtained: for instance an integrated index of impact, to measure scouring on buildings by multiple waves in order to better appraise potential damages to buildings. Furthermore, early warning systems based on precomputed solutions, for example, Gusman et al. (2014), could be substantially upgraded using such an approach. The use of the fast emulator could help issue more precise warnings, accompanied by uncertainties arising from the large uncertain initial conditions inferred from early estimates of the earthquake characteristics. Further work includes a better description of the earthquake characteristics, with more parameters. Such advances will present further computational and geophysical challenges in order to adequately quantify uncertainties in tsunami risk for the Cascadia region. Our approach could also be utilized in other tsunami-prone regions of the world.

Acknowledgments. The work presented here made use of the Emerald GPU supercomputer, which was made available by the Center for Innovation in High

Performance Computing. The Center is formed by the universities of Oxford, Southampton, and Bristol and University College London in partnership with the STFC Rutherford Appleton Laboratory, through EPSRC funding EP/K000144/1.

SUPPLEMENTARY MATERIAL

Supplement A: Two movies of seabed deformations (DOI: 10.1214/18-AOAS1142SUPP; .zip). These two seabed deformations correspond to cases 5 and 13 in the Latin Hypercube Design of 40 combinations of the three input characteristics.

REFERENCES

ADAMS, J. (1990). Paleoseismicity of the Cascadia subduction zone: Evidence from turbidites off the Oregon–Washington margin. *Tectonics* **9** 569–583.

ATWATER, F. B., NELSON, R. A., CLAGUE, J. J., CARVER, A. G., YAMAGUCHI, K. D., BOBROWSKY, T. P., BOURGEOIS, J., DARIENZO, E. M., GRANT, C. W., HEMPHILL-HALEY, E., KELSEY, M. H., JACOBY, C. G., NISHENKO, P. S., PALMER, P. S., PETERSON, D. C. and REINHART, M. A. (1995). Summary of coastal geologic evidence for past great earthquakes at the Cascadia subduction zone. *Earthq. Spectra* **11** 1–18.

BAYARRI, M. J., BERGER, J. O., CAFEO, J., GARCIA-DONATO, G., LIU, F., PALOMO, J., PARTHASARATHY, R. J., PAULO, R., SACKS, J. and WALSH, D. (2007). Computer model validation with functional output. *Ann. Statist.* **35** 1874–1906. MR2363956

BECK, J. and GUILLAS, S. (2016). Sequential design with mutual information for computer experiments (MICE): Emulation of a tsunami model. *SIAM/ASA J. Uncertain. Quantificat.* **4** 739–766. MR3507556

BEHRENS, J. and DIAS, F. (2015). New computational methods in tsunami science. *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **373** 20140382. MR3405409

BERNARD, E. and TITOV, V. (2015). Evolution of tsunami warning systems and products. *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **373** 20140371.

BILIONIS, I., ZABARAS, N., KONOMI, B. A. and LIN, G. (2013). Multi-output separable Gaussian process: Towards an efficient, fully Bayesian paradigm for uncertainty quantification. *J. Comput. Phys.* **241** 212–239. MR3647406

BRICKER, J. D. and NAKAYAMA, A. (2014). Contribution of trapped air, deck superelevation, and nearby structures to bridge deck failure during a tsunami. *J. Hydraul. Eng.* **140** 05014002.

CASTRO, P. E., LAWTON, W. H. and SYLVESTRE, E. A. (1986). Principal modes of variation for processes with continuous sample curves. *Technometrics* **28** 329–337.

CHANG, W., APPLEGATE, P. J., HARAN, M. and KELLER, K. (2014). Probabilistic calibration of a Greenland Ice Sheet model using spatially resolved synthetic observations: Toward projections of ice mass loss with uncertainties. *Geosci. Model Dev.* **7** 1933–1943.

DAY, S. and FEARNLEY, C. (2015). A classification of mitigation strategies for natural hazards: Implications for the understanding of interactions between mitigation strategies. *Nat. Hazards* **79** 1219–1238.

DUTYKH, D., PONCET, R. and DIAS, F. (2011). The VOLNA code for the numerical modeling of tsunami waves: Generation, propagation and inundation. *Eur. J. Mech. B Fluids* **30** 598–615. MR2906226

GILES, M. B., MUDALIGE, G. R., SHARIF, Z., MARKALL, G. and KELLY, P. H. (2011). Performance analysis of the OP2 framework on many-core architectures. *ACM SIGMETRICS Perform. Eval. Rev.* **38** 9–15.

GOLDFINGER, C., NELSON, C. H., JOHNSON, J. E., MOREY, A. E., GUTIERREZ-PASTOR, J., KARABANOV, E., ERIKSSON, A. T., GRACIA, E., DUNHILL, G., PATTON, J., ENKIN, R., DALLIMORE, A. and VALLIER, T. (2012). Turbidite event history: Methods and implications for Holocene paleoseismicity of the Cascadia Subduction Zone. U.S. Geol. Surv. Professional Paper 1661.

GONZÁLEZ, F. I., GEIST, E. L., JAFFE, B., KÂNOĞLU, U., MOFJELD, H., SYNOLAKIS, C. E., TITOV, V. V., ARCAS, D., BELLOMO, D., CARLTON, D. et al. (2009). Probabilistic tsunami hazard assessment at seaside, Oregon, for near- and far-field seismic sources. *J. Geophys. Res., Oceans* **114** C11023.

GUILLAS, S., SARRI, A., DAY, S. J., LIU, X. and DIAS, F. (2018). Supplement to "Functional emulation of high resolution tsunami modelling over Cascadia." DOI:10.1214/18-AOAS1142SUPP.

GUSMAN, A. R., TANIOKA, Y., MACINNES, B. T. and TSUSHIMA, H. (2014). A methodology for near-field tsunami inundation forecasting: Application to the 2011 Tohoku tsunami. *J. Geophys. Res., Solid Earth* **119** 8186–8206.

HAWKES, A. D., HORTON, B. P., NELSON, A. R., VANE, C. H. and SAWAI, Y. (2011). Coastal subsidence in Oregon, USA, during the giant Cascadia earthquake of AD 1700. *Quat. Sci. Rev.* **30** 364–376.

HEIDARZADEH, M., PIROOZ, M. D. and ZAKER, N. H. (2009). Modeling the near-field effects of the worst-case tsunami in the Makran subduction zone. *Ocean Eng.* **36** 368–376.

HIGDON, D., GATTIKER, J., WILLIAMS, B. and RIGHTLEY, M. (2008). Computer model calibration using high-dimensional output. *J. Amer. Statist. Assoc.* **103** 570–583. MR2523994

HUNG, Y., JOSEPH, V. R. and MELKOTE, S. N. (2015). Analysis of computer experiments with functional response. *Technometrics* **57** 35–44. MR3318347

HYNDMAN, R. D. (2013). Downdip landward limit of Cascadia great earthquake rupture. *J. Geophys. Res., Solid Earth* **118** 5530–5549.

JOLLIFFE, I. T. (2002). *Principal Component Analysis*, 2nd ed. Springer, New York. MR2036084

JOSEPH, V. R., HUNG, Y. and SUDJIANTO, A. (2008). Blind kriging: A new method for developing metamodels. *J. Mech. Des.* **130** 031102.

KAUFMAN, C. G., BINGHAM, D., HABIB, S., HEITMANN, K. and FRIEMAN, J. A. (2011). Efficient emulators of computer experiments using compactly supported correlation functions, with an application to cosmology. *Ann. Appl. Stat.* **5** 2470–2492. MR2907123

KLEIBER, W., KATZ, R. W. and RAJAGOPALAN, B. (2013). Daily minimum and maximum temperature simulation over complex terrain. *Ann. Appl. Stat.* **7** 588–612. MR3086432

KNEIP, A. and RAMSAY, J. O. (2008). Combining registration and fitting for functional models. *J. Amer. Statist. Assoc.* **103** 1155–1165. MR2528838

LAY, T., KANAMORI, H., AMMON, C. J., NETTLES, M., WARD, S. N., ASTER, R. C., BECK, S. L., BILEK, S. L., BRUDZINSKI, M. R., BUTLER, R. et al. (2005). The great Sumatra-Andaman earthquake of 26 December 2004. *Science* **308** 1127–1133.

LAY, T., AMMON, C. J., KANAMORI, H., KOPER, K. D., SUFRI, O. and HUTKO, A. R. (2010). Teleseismic inversion for rupture process of the 27 February 2010 Chile (Mw 8.8) earthquake. *Geophys. Res. Lett.* **37** L13301.

LEONARD, L. J., ROGERS, G. C. and MAZZOTTI, S. (2014). Tsunami hazard assessment of Canada. *Nat. Hazards* **70** 237–274.

LEONARD, L. J., CURRIE, C. A., MAZZOTTI, S. and HYNDMAN, R. D. (2010). Rupture area and displacement of past Cascadia great earthquakes from coastal coseismic subsidence. *Geol. Soc. Am. Bull.* **122** 2079–2096.

LIU, X. and GUILLAS, S. (2017). Dimension reduction for Gaussian process emulation: An application to the influence of bathymetry on tsunami heights. *SIAM/ASA J. Uncertain. Quantificat.* **5** 787–812. MR3686816

LÓPEZ-PINTADO, S. and ROMO, J. (2009). On the concept of depth for functional data. *J. Amer. Statist. Assoc.* **104** 718–734. MR2541590

LUDWIN, R. S., DENNIS, R., CARVER, D., MCMILLAN, A. D., LOSEY, R., CLAGUE, J., JONIENTZ-TRISLER, C., BOWECHOP, J., WRAY, J. and JAMES, K. (2005). Dating the 1700 Cascadia earthquake: Great coastal earthquakes in native stories. *Seismol. Res. Lett.* **76** 140–148.

MASTERLARK, T. and HUGHES, K. L. (2008). Next generation of deformation models for the 2004 M9 Sumatra-Andaman earthquake. *Geophys. Res. Lett.* **35** L19310.

MAZZOTTI, S., DRAGERT, H., HENTON, J., SCHMIDT, M., HYNDMAN, R., JAMES, T., LU, Y. and CRAYMER, M. (2003). Current tectonics of northern Cascadia from a decade of GPS measurements. *J. Geophys. Res., Solid Earth* **108** 2554.

MOORE, G. F., BANGS, N. L., TAIRA, A., KURAMOTO, S., PANGBORN, E. and TOBIN, H. J. (2007). Three-dimensional splay fault geometry and implications for tsunami generation. *Science* **318** 1128–1131.

MORRIS, M. D. (2012). Gaussian surrogates for computer models with time-varying inputs and outputs. *Technometrics* **54** 42–50. MR2904734

MUDALIGE, G. R., GILES, M. B., REGULY, I., BERTOLLI, C. and KELLY, P. H. J. (2012). OP2: An active library framework for solving unstructured mesh-based applications on multi-core and many-core architectures. In 2012 *Innovative Parallel Computing* (*InPar*) 1–12. IEEE, Los Alamitos, CA.

MUDALIGE, G. R., GILES, M. B., THIYAGALINGAM, J., REGULY, I. Z., BERTOLLI, C., KELLY, P. H. and TREFETHEN, A. E. (2013). Design and initial performance of a high-level unstructured mesh framework on heterogeneous parallel systems. *Parallel Comput.* **39** 669–692.

NELSON, A. R., SAWAI, Y., JENNINGS, A. E., BRADLEY, L.-A., GERSON, L., SHERROD, B. L., SABEAN, J. and HORTON, B. P. (2008). Great-earthquake paleogeodesy and tsunamis of the past 2000 years at Alsea Bay, central Oregon coast, USA. *Quat. Sci. Rev.* **27** 747–768.

OAKLEY, J. E. and O'HAGAN, A. (2004). Probabilistic sensitivity analysis of complex models: A Bayesian approach. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **66** 751–769. MR2088780

OKADA, Y. (1992). Internal deformation due to shear and tensile faults in a half-space. *Bull. Seismol. Soc. Am.* **82** 1018–1040.

OKAL, E. A. (2015). The quest for wisdom: Lessons from 17 tsunamis, 2004–2014. *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **373** 20140370.

RAMSAY, J., HOOKER, G. and GRAVES, S. (2009). *Functional Data Analysis with R and MATLAB*. Springer, Dordrecht.

RAMSAY, J. O. and SILVERMAN, B. W. (2005). *Functional Data Analysis*, 2nd ed. Springer, New York. MR2168993

ROUGIER, J. (2008). Efficient emulators for multivariate deterministic functions. *J. Comput. Graph. Statist.* **17** 827–843. MR2649069

ROUGIER, J., GUILLAS, S., MAUTE, A. and RICHMOND, A. D. (2009). Expert knowledge and multivariate emulation: The thermosphere-ionosphere electrodynamics general circulation model (TIE-GCM). *Technometrics* **51** 414–424. MR2756477

SALMANIDOU, D. M., GUILLAS, S., GEORGIOPOULOU, A. and DIAS, F. (2017). Statistical emulation of landslide-induced tsunamis at the Rockall Bank, NE Atlantic. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **473** 20170026. MR3650585

SARRI, A., GUILLAS, S. and DIAS, F. (2012). Statistical emulation of a tsunami model for sensitivity analysis and uncertainty quantification. *Nat. Hazards Earth Syst. Sci.* **12** 2003–2018.

SATAKE, K., WANG, K. and ATWATER, B. F. (2003). Fault slip and seismic moment of the 1700 Cascadia earthquake inferred from Japanese tsunami descriptions. *J. Geophys. Res., Solid Earth* **108** 2535.

SATAKE, K., SHIMAZAKI, K., TSUJI, Y. and UEDA, K. (1996). Time and size of a giant earthquake in Cascadia inferred from Japanese tsunami records of January 1700. *Nature* **379** 246–249.

SATAKE, K., FUJII, Y., HARADA, T. and NAMEGAYA, Y. (2013). Time and space distribution of coseismic slip of the 2011 tohoku earthquake as inferred from tsunami waveform data. *Bull. Seismol. Soc. Am.* **103** 1473–1492.

SILVERMAN, B. W. (1996). Smoothed functional principal components analysis by choice of norm. *Ann. Statist.* **24** 1–24. MR1389877

SIMONS, M., MINSON, S. E., SLADEN, A. et al. (2011). The 2011 magnitude 9.0 Tohoku-Oki earthquake: Mosaicking the megathrust from seconds to centuries. *Science* **332** 1421–1425.

SPILLER, E. T., BAYARRI, M. J., BERGER, J. O., CALDER, E. S., PATRA, A. K., PITMAN, E. B. and WOLPERT, R. L. (2014). Automating emulator construction for geophysical hazard maps. *SIAM/ASA J. Uncertain. Quantificat.* **2** 126–152. MR3283903

WANG, K., MULDER, T., ROGERS, G. C. and HYNDMAN, R. D. (1995). Case for very low coupling stress on the Cascadia Subduction Fault. *J. Geophys. Res., Solid Earth* **100** 12907–12918.

WANG, P.-L., ENGELHART, S. E., WANG, K., HAWKES, A. D., HORTON, B. P., NELSON, A. R. and WITTER, R. C. (2013). Heterogeneous rupture in the great Cascadia earthquake of 1700 inferred from coastal subsidence estimates. *J. Geophys. Res., Solid Earth* **118** 2460–2473.

WELLS, D. L. and COPPERSMITH, K. J. (1994). New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement. *Bull. Seismol. Soc. Am.* **84** 974–1002.

WITTER, R. C., ZHANG, Y. J., WANG, K., PRIEST, G. R., GOLDFINGER, C., STIMELY, L., ENGLISH, J. T. and FERRO, P. A. (2013). Simulated tsunami inundation for a range of Cascadia megathrust earthquake scenarios at Bandon, Oregon, USA. *Geosphere* **9** 1783–1803.

S. GUILLAS
A. SARRI
S. J. DAY
X. LIU
UNIVERSITY COLLEGE LONDON
GOWER STREET
LONDON, WC1E 6BT
UNITED KINGDOM
E-MAIL: s.guillas@ucl.ac.uk

F. DIAS
SCHOOL OF MATHEMATICS AND STATISTICS
UNIVERSITY COLLEGE DUBLIN
SCIENCE CENTRE—NORTH
BELFIELD, DUBLIN 4
IRELAND