

Adjoint-based sensitivity analysis for a numerical storm surge model

Simon C. Warder^{a,*}, Kevin J. Horsburgh^b, Matthew D. Piggott^a

^a*Department of Earth Science and Engineering, Imperial College London, London, SW7 2AZ, UK*

^b*National Oceanography Centre, Liverpool, L3 5DA, UK*

Abstract

Numerical storm surge models are essential to forecasting coastal flood hazard and informing the design of coastal defences. However, such models rely on a variety of inputs, many of which carry uncertainty, and an awareness and understanding of the sensitivity of the model outputs with respect to those uncertain inputs is necessary when interpreting model results. Here, we use an unstructured-mesh numerical coastal ocean model, *Thetis*, and its adjoint, to perform a sensitivity analysis for a hindcast of the 5th/6th December 2013 North Sea surge event, with respect to the bottom friction coefficient, bathymetry and wind stress forcing. The results reveal spatial and temporal patterns of sensitivity, providing physical insight into the mechanisms of surge generation and propagation. For example, the sensitivity of the skew surge to the bathymetry reveals the protective effect of a sand bank off the UK east coast. The results can also be used to propagate uncertainties through the numerical model; based on estimates of model input uncertainties, we estimate that modelled skew surges carry uncertainties of around 5 cm and 15 cm due to bathymetry and bottom friction, respectively. While these uncertainties are small compared with the typical spread in an ensemble storm surge forecast due to uncertain meteorological inputs, the adjoint-derived model sensitivities can nevertheless be used to inform future model calibration and data acquisition efforts in order to reduce uncertainty. Our results demonstrate the power of adjoint methods to gain relevant insight into a storm surge model, providing information complementary to traditional ensemble uncertainty quantification methods.

Keywords: Storm surge, Adjoint, Sensitivity analysis, Uncertainty quantification, Unstructured mesh, Finite element method

*Corresponding author

Email addresses: s.warder15@imperial.ac.uk (Simon C. Warder), kevinh@noc.ac.uk (Kevin J. Horsburgh), m.d.piggott@imperial.ac.uk (Matthew D. Piggott)

1. Introduction

Storm surge poses a significant hazard for coastal communities worldwide. Allowing for investment in adaptation measures (e.g. rising flood defences), global flood losses in 136 of the world’s largest coastal cities have recently been estimated to rise from US\$6 bn per year in 2005 to US\$60-63 bn per year in 2050 (Hallegatte et al., 2013). Globally, the increase in extreme sea levels (Stocker et al., 2013) will result in critical flood defence thresholds being reached more frequently and therefore the risk of flooding will increase. The UK is vulnerable to storm surges, particularly along its North Sea coast; a large number of severe storms have impacted the UK in the last century (Haigh et al., 2016), with the two most severe of those events occurring in the North Sea in 1953 and 2013. The approximate economic impacts of the coastal flooding resulting from these events (for year 2014) were £1.2 bn and £0.25 bn respectively; the impact of the latter event was reduced through mitigation action taken after the 1953 event (Wadey et al., 2015). With continued development of the coastal zone in flood risk areas (ASC, 2014), the role of storm surge modelling remains vital.

Essential to the intelligent application of any storm surge model is an understanding of the model’s sensitivity to its uncertain inputs. In a forecast scenario, the greatest model uncertainty arises from the meteorological forcing, namely the surface stress due to wind, and the atmospheric pressure gradient. For this reason, it is common to employ ensemble methods for uncertainty quantification, whereby the surge model is run multiple times, with each run using a different sample from the uncertain distribution of meteorological inputs (Flowerdew et al., 2010). While such ensemble methods provide a practical approach to uncertainty quantification within an operational forecast framework, they provide little insight into the patterns (in space and/or time) of the underlying model sensitivity, and they depend on the choice of meteorological ensemble.

An alternative approach to sensitivity analysis is provided by adjoint methods. In the context of numerical modelling, adjoint methods are used to efficiently compute gradients of model outputs with respect to model inputs, which can in principle vary in both space and time. Such methods have been used within a meteorological context since the 1980s (e.g. Hall et al. (1982)), and have a variety of applications within the field of coastal ocean modelling. Adjoint-derived sensitivities to model inputs can be used for gaining physical insight into a modelled system (e.g. Losch and Heimbach (2007), Massmann (2010), Verdy et al. (2014), Nowak (2015), Villaret et al. (2016)), or can be used within frameworks for model calibration, data assimilation and parameter estimation (e.g. Lardner et al. (1993), Canizares et al. (1998), Heemink et al. (2002), Lu and Zhang (2006), Zhang et al. (2011), Li et al. (2013), Chen et al. (2014)). Adjoint methods have previously been applied to the analysis of storm surge model sensitivity to wind stress (Wilson et al., 2013, Warder et al., 2019), and this paper represents an extension to these works.

Here, we apply a numerical coastal ocean model, *Thetis*, and its adjoint, to perform a storm surge sensitivity analysis with respect to multiple model inputs, namely the bottom friction coefficient, bathymetry and wind stress. We use the resulting sensitivities to gain physical insight into surge generation and propagation

37 in the North Sea, and to estimate and compare the uncertainty in surge model outputs arising from each
 38 of these inputs, and at different locations in the model domain. We first introduce the numerical model in
 39 section 2, and perform a brief model calibration in section 3. The adjoint approach to sensitivity analysis is
 40 described in section 4, and sensitivity analysis results are presented in section 5, using the extreme December
 41 2013 storm surge event as a case study. The results of the sensitivity analysis are discussed in section 6, and
 42 conclusions are made in section 7.

43 2. Forward numerical model

44 Within this work, we model storm surges using *Thetis*, an unstructured-mesh finite element coastal ocean
 45 flow solver (Kärnä et al., 2018) implemented within the Firedrake finite element code generation framework
 46 (Rathgeber et al., 2016). We use Thetis in its two-dimensional configuration (Vouriot et al., 2019), which
 47 solves the shallow water equations (SWEs) in non-conservative form, given by

$$\begin{aligned} \frac{\partial \eta}{\partial t} + \nabla \cdot (H\mathbf{u}) &= 0, \\ \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \mathbf{F}_C + g\nabla \eta + \nabla \left(\frac{p_a}{\rho} \right) &= \frac{\boldsymbol{\tau}_s - \boldsymbol{\tau}_b}{\rho H} + \nabla \cdot (\nu_h(\nabla \mathbf{u} + \nabla \mathbf{u}^T)), \end{aligned} \quad (1)$$

48 where η is the free surface height, H is the water depth given by $H = \eta + h$ where h is the bathymetry
 49 (measured positive downwards), \mathbf{u} is the two-dimensional depth-averaged velocity vector, \mathbf{F}_C is the Coriolis
 50 force, g is the acceleration due to gravity, ρ is the water density, p_a is the atmospheric pressure at the free
 51 surface, $\boldsymbol{\tau}_b$ is the bottom stress due to friction with the sea bed, $\boldsymbol{\tau}_s$ is the surface stress due to wind, and ν_h
 52 is the kinematic viscosity. All variables are in SI units.

53 The bathymetry h is taken from the GEBCO 2014 dataset, which has a resolution of 30 arc-seconds (or
 54 approximately 1 km), and is linearly interpolated onto the model mesh. A minimum depth of 10 m is then
 55 applied to avoid the need for wetting and drying, since we do not resolve the spatial scales of inundation
 56 within this study.

57 The bottom friction is parameterised via Manning’s n formulation, such that

$$\frac{\boldsymbol{\tau}_b}{\rho} = \frac{gn^2}{H^{\frac{1}{3}}} |\mathbf{u}| \mathbf{u}, \quad (2)$$

58 where n is the Manning coefficient. We assume that the friction coefficient is spatially uniform and constant
 59 in time; this assumption is consistent with CS3X, the UK operational surge forecast model at the time of
 60 this event (Flowerdew et al., 2013). The Manning coefficient n is determined by a preliminary calibration
 61 exercise, as described in section 3.

62 Storm surge forcing is included via p_a and $\boldsymbol{\tau}_s$, which are the atmospheric pressure and surface stress due to
 63 wind, respectively. Within this work, we focus on the surge event of 5th/6th December 2013. Meteorological
 64 hindcast data for this event were provided by the National Oceanography Centre (personal communication
 65 2018) at a spatial resolution of $1/9^\circ$ latitude by $1/6^\circ$ longitude, and a temporal resolution of 1 hour. This
 66 forcing data is linearly interpolated onto our model mesh and between timesteps for use within Thetis. We

67 use a Charnock parameterisation similar to Williams and Flather (2000) to relate τ_s to the wind velocity,
 68 via the system of equations

$$\begin{aligned}\tau_s &= \rho_{\text{air}} |\mathbf{W}_*| \mathbf{W}_*, \\ \mathbf{W} &= \mathbf{W}_* \frac{1}{\kappa} \log \frac{z}{z_0}, \\ z_0 &= \frac{\alpha |\mathbf{W}_*|^2}{g},\end{aligned}\tag{3}$$

69 where ρ_{air} is the density of air, \mathbf{W}_* is the friction velocity, \mathbf{W} the wind velocity at height z above the free
 70 surface, assuming neutral atmospheric stratification, κ is the von Kármán constant, taken to be 0.4, z_0 the
 71 surface roughness, and α the Charnock parameter, which we select via a calibration exercise as described
 72 in section 3. This Charnock parameter is assumed within this work to be spatially uniform and constant
 73 in time, neglecting any variation due to surface waves. This formulation is a common choice within storm
 74 surge modelling, and is consistent with that of the CS3X model (Williams and Flather, 2000, Flowerdew
 75 et al., 2010, 2013). The system of equations (3) is solved by a simple iterative method.

76 Tidal forcing is included within the model by applying a Dirichlet boundary condition for the free surface
 77 height on the open ocean boundaries, generated from eight harmonic constituents from the TPXO database
 78 (Egbert and Erofeeva, 2002) (M2, S2, N2, K2, Q1, O1, P1, K1). This boundary condition is further modified
 79 by a correction calculated from the inverse barometer effect, which is applied to approximate surge generated
 80 externally to the model domain.

81 The governing equations (1) are solved on an unstructured mesh using a $P_1^{\text{DG}}-P_1^{\text{DG}}$ finite element pair, us-
 82 ing approximate Riemann fluxes at element interfaces as described in Kärnä et al. (2018). This discretisation
 83 has been shown to be well suited for shallow water problems (Comblen et al., 2010). We use a Crank-Nicolson
 84 timestepping scheme with a timestep of 100 s. The mesh used within this work is shown in figure 1, in addi-
 85 tion to a close-up of a key region of the domain. The mesh was generated using the Python package *qmesh*
 86 (version 1.0.1) (Avdis et al., 2018), which interfaces the mesh generator *Gmsh* (version 2.10.1) (Geuzaine
 87 and Remacle, 2009). The mesh uses the UTM31 coordinate system, and its resolution varies from 3 km at
 88 the coastline to 25 km in the open ocean, resulting in a total of 23,120 triangular elements. This coastline
 89 resolution is finer than that of the CS3X model (which has approximately 12 km resolution) (Flowerdew
 90 et al., 2013), while also benefitting from the alignment of element edges with the coastline, a key advantage
 91 of unstructured mesh methods (Pain et al., 2005). In open ocean regions, 25 km resolution is assumed to
 92 sufficiently capture the dynamics, since the tidal wavelength is $\mathcal{O}(1000\text{km})$. Each model run for this study
 93 was performed in parallel using 16 cores, on a machine with 124 GB memory, and each 30-day forward model
 94 run took a wall-clock time of approximately 7.5 hours.

95 3. Model calibration

96 We first calibrate the model with respect to the Manning coefficient n , based on a tide-only simulation.
 97 After a spin-up period of 10 days, the model is run in tide-only mode for one month, and a harmonic analysis

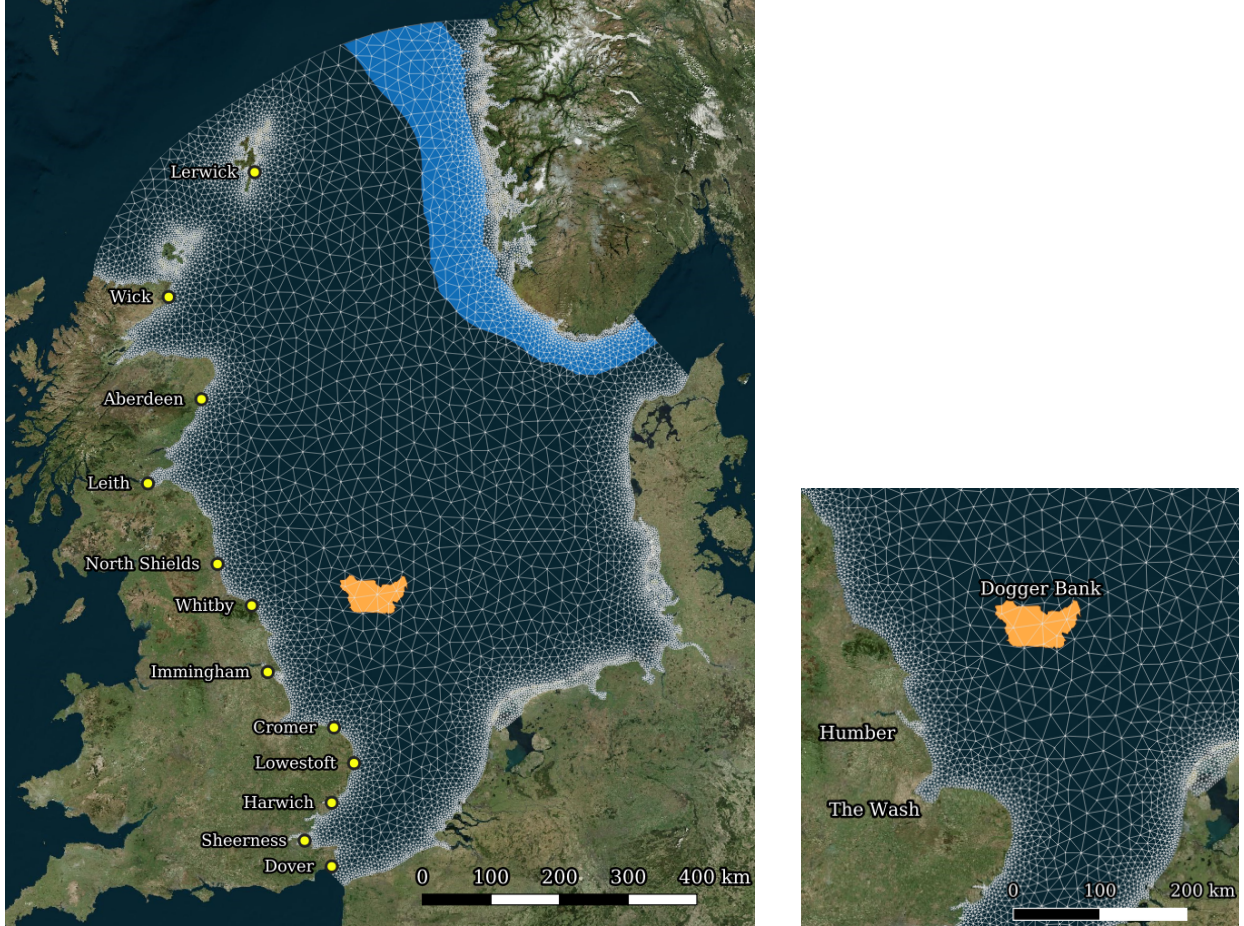


Figure 1: Left: Mesh used for all simulations within this work, consisting of 23,120 triangular elements. Tide gauge locations are shown for the east coast of the UK mainland. Two bathymetric features are highlighted for later reference: Dogger Bank (orange, extracted from 20 m bathymetry contour) and the Norwegian Trench (blue, extracted from 200 m bathymetry contour). Right: Close-up of a key region of the model domain, featuring the Humber, the Wash, and Dogger Bank.

98 performed at the 12 tide gauge stations within the model domain where quality controlled data is available
 99 from the British Oceanographic Data Centre (BODC) (see figure 1). We perform this harmonic analysis
 100 based on the same eight harmonic constituents as the tidal boundary condition, using the Python package
 101 *uptide* (Kramer et al., 2020). The model-observation error is computed via the combined root mean squared
 102 error (RMSE) of the amplitudes of the eight harmonic constituents C by

$$\text{RMSE} = \left(\frac{1}{8} \sum_C \frac{1}{12} \sum_{i=1}^{12} (A_{C,i} - \hat{A}_{C,i})^2 \right)^{\frac{1}{2}}, \quad (4)$$

103 where $A_{C,i}$ and $\hat{A}_{C,i}$ are the modelled and observed amplitudes of the harmonic constituent C at tide gauge
 104 location i , respectively. The model was run as described above, for values of the Manning coefficient n from
 105 0.015 $\text{s m}^{-1/3}$ to 0.04 $\text{s m}^{-1/3}$ in steps of 0.0025 $\text{s m}^{-1/3}$. The results are shown in figure 2; the smallest
 106 value for the RMSE was 5.7 cm, achieved with $n = 0.025 \text{ s m}^{-1/3}$. This value is used for the remainder of
 107 this paper. We note that the use of a spatially uniform value for the friction coefficient may restrict the




Figure 2: Calibration of the tidal model with respect to the Manning coefficient, n . The minimum RMSE is achieved using $n = 0.025 \text{ s m}^{-1/3}$.

108 predictive capability of the model, but the RMSE attained with a uniform friction parameter is considered
 109 adequate, and more sophisticated model calibration is outside the scope of this work. It is assumed that the
 110 model sensitivities to the friction coefficient are not strongly dependent on the value of the coefficient itself.

111 In order to select an appropriate value for the Charnock parameter α , the surge model was run for the
 112 December 2013 event using varying values of α . For these simulations, the model is first spun up (in tide-only
 113 mode) for 10 days, prior to the wind and atmospheric pressure forcing terms being switched on approximately
 114 10 days before the peak storm tide occurs. A comparison with observations for this event is made based
 115 on the modelled and observed surge residuals at the BODC tide gauge locations. As shown in figure 3, the
 116 surge residual is defined as the difference between the storm tide (i.e. the total sea surface height due to tidal
 117 and meteorological forcing) and the astronomical tide (i.e. the sea surface height which would be expected
 118 in the absence of meteorological forcing). For the tide gauge observations, the astronomical component is
 119 computed based on a harmonic analysis of long term tide gauge data, and this harmonic part is subtracted
 120 from the observed sea surface heights to obtain the residual. To compute the modelled residual, the model
 121 is simply run in both full surge (tidally and meteorologically forced) and tide-only modes, and the surface
 122 elevations subtracted. For the purposes of calibrating the Charnock parameter, the model-observation error
 123 is computed by a simple root mean squared error of the residual timeseries, over a two day period capturing
 124 the peak storm tide, at the eight BODC tide gauges within the domain at which the surge was significant,
 125 and which recorded a sufficiently complete timeseries surface elevation record during the event. The surge
 126 model was run as described, for values of the Charnock parameter α from 0.01 to 0.03, in steps of 0.002.
 127 The results are shown in figure 4; the smallest value for the residual RMSE was 15.9 cm, obtained using
 128 $\alpha = 0.028$, which is a value consistent with the literature (Williams and Flather, 2000, Brown and Wolf,
 129 2009).

130 Using these calibrated parameters, a good agreement is obtained between modelled and observed surge

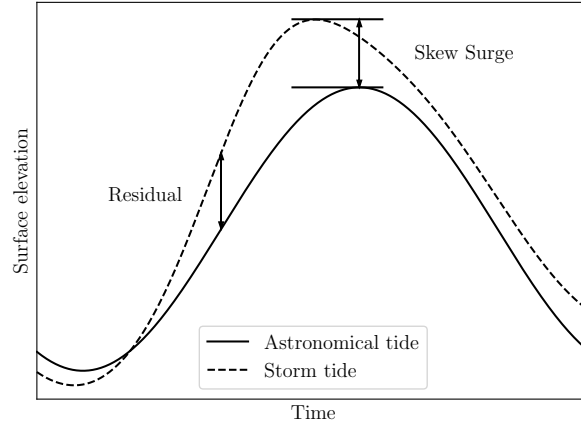


Figure 3: Schematic diagram of residual and skew surge definitions. The astronomical tide is the sea surface level which would be observed due to astronomical tidal forcing only. The storm tide is the sea surface level induced by the combination of the astronomical tide and meteorological conditions.

131 residuals for this event, as shown in figure 5 for the three tide gauge locations selected for the sensitivity
 132 analysis study. The smoothness of the model outputs compared with the observations arises from several
 133 factors, including the choice of mesh, discretisation (e.g. choice of finite element pair), timestepping, the
 134 representation within the model of the bathymetry and coastlines, and limitations in other model inputs.
 135 While adjustments to the model setup, e.g. the use of higher order mesh elements, might improve its
 136 predictive capability, the attained model-observation agreement is considered sufficient for the purposes of
 137 this study.

138 As an additional experiment, we tested alternative model meshes with (i) uniform 12 km resolution, and
 139 (ii) coastline resolution of 1.5 km and open-ocean resolution of 15 km. The residual RMSEs computed with
 140 these meshes were both within 1 cm of the RMSEs obtained using the final mesh we selected. This suggests
 141 that the selected mesh resolution does not limit our modelling accuracy, and that the model-observation
 142 misfit is dominated by other factors.




Figure 4: Calibration of the tidal model with respect to the Charnock parameter, α . The minimum RMSE is achieved using $\alpha = 0.028$.

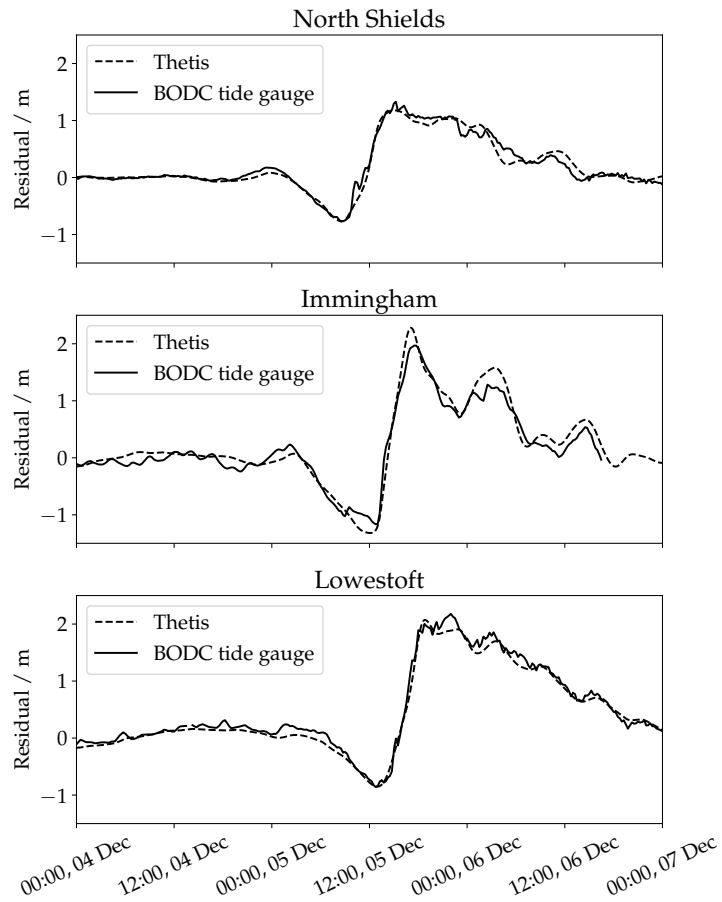


Figure 5: Comparison of modelled and observed surge residuals for the December 2013 event, at three selected tide gauge locations. The RMSEs between the modelled and observed residuals at North Shields, Immingham and Lowestoft are 11.1, 19.6 and 11.3 cm, respectively.

143 **4. Methods**

144 *4.1. The adjoint method*

145 In this section we briefly describe the adjoint method (based largely on Funke (2012)), and for further
 146 detail the reader is referred to similar works in the literature (e.g. Wilson et al. (2013), Verdy et al. (2014))
 147 and previous studies utilising the adjoint mode of Thetis (e.g. Warder et al. (2019), Goss et al. (2020)).

148 For compactness, we write our system of PDEs in the general form

$$F(u, m) = 0, \tag{5}$$

149 where F is the PDE operator (representing the shallow water equations (1)), u is the model state variable
 150 (representing η and \mathbf{u}) and m represents the input parameters (bottom friction n , bathymetry h and wind
 151 stress $\boldsymbol{\tau}_s$). A functional of interest is given by a scalar function $J(u)$. The purpose of the adjoint method is to
 152 efficiently compute the derivative of J with respect to the input parameters m , i.e. $\frac{dJ}{dm}$. It is straightforward
 153 to show that

$$\frac{dJ}{dm} = -\frac{\partial J}{\partial u} \frac{\partial F^{-1}}{\partial u} \frac{\partial F}{\partial m}. \tag{6}$$

154 When solved in the forward time direction, this represents the so-called tangent linear approach to the
 155 evaluation of the gradient $\frac{dJ}{dm}$. However, this approach is inefficient when there are a small number of
 156 functionals J , and a large number of parameters m , as is the case within this work. We therefore take the
 157 so-called adjoint approach. Equation (6) can be expressed as

$$\frac{dJ}{dm} = -\lambda^* \frac{\partial F}{\partial m}, \tag{7}$$

158 where λ is the adjoint variable, defined as the solution to the adjoint equation

$$\frac{\partial F^*}{\partial u} \lambda = \frac{\partial J^*}{\partial u}, \tag{8}$$

159 where the asterisk denotes the adjoint operation (that is, the conjugate transpose). Note that this is a
 160 PDE; the adjoint variable λ is space- and time-dependent, and the adjoint equation is solved in the reverse
 161 time direction (supposing that the functional J depends only on the final state of the model variables, it
 162 is intuitive to interpret the right hand side of this equation as an initial condition, which is propagated
 163 backwards in time by the adjoint equation). Since the right-hand-side of equation (7) is a time-varying
 164 quantity, it represents the instantaneous influence of the parameter m on the functional J . In the case where
 165 the inputs m are not time-varying (i.e. for m representing the bottom friction coefficient or bathymetry),
 166 the right-hand-side of this equation should be integrated over all time (although in practice here we integrate
 167 over a period of 10 days).

168 Note that the adjoint equation is always linear, and that the parameter m does not appear in the equation.
 169 The computational cost of the numerical solution of the adjoint equation is therefore reasonably low, and
 170 does not depend on the number of parameters.



Figure 6: Selected target locations for sensitivity analysis. Three tide gauge locations are selected, along with an integrated measure along the indicated section of coastline.

171 For an application to numerical models, u and m in the above exposition should be interpreted as vectors
 172 representing the model state and input parameters, respectively. Specifically, m corresponds to a vector of
 173 model inputs, whose elements correspond to the values of the input at each mesh node (and, in the case of
 174 the wind stress, at each time step). For models, such as Thetis, which are implemented within the Firedrake
 175 framework, the adjoint model can be generated algorithmically via the Python package *pyadjoint* (Mitusch
 176 et al., 2019, Farrell et al., 2013), removing the need to derive the adjoint equations by hand and implement
 177 their numerical solution.

178 For the model setup within this work, runs of Thetis in its adjoint mode were found to require approxi-
 179 mately 2.4 times the wall-clock time of the forward mode. We emphasise that, since this computational cost
 180 is independent of the number of parameters m , the adjoint approach is especially powerful when used to
 181 calculate the derivative of a small number of model outputs with respect to a large number of inputs. Since
 182 the model inputs are defined at each mesh node (and at each time step in the case of wind stress), within
 183 this work the vector of input parameters m contains $\mathcal{O}(10^6)$ elements. The adjoint method is therefore the
 184 only feasible approach to computing the sensitivity $\frac{dJ}{dm}$.

185 4.2. Sensitivity analysis and uncertainty quantification

186 Taking the December 2013 surge event, we use four definitions for the functional J . As indicated in
 187 figure 6, these correspond to the skew surges at three tide gauge locations (North Shields, Immingham
 188 and Lowestoft), and the mean skew surge along a section of coastline (measuring approximately 400 km,
 189 from Bridlington in the north to Great Yarmouth in the south, and including the Humber Estuary and the
 190 Wash). The skew surge is defined as the difference between the peak storm tide surface elevation and the
 191 peak astronomical (tidal) surface elevation, as shown in figure 3. The choice of skew surge as the selected

192 model output at each target location is motivated by Williams et al. (2016); the skew surge constitutes a
 193 more meaningful measure of surge severity than the surge residual. A functional defined as the skew surge
 194 at a location \mathbf{x}_0 can be broken down as

$$J_{\text{skew surge}} = J_{\text{peak storm tide}} - J_{\text{peak astronomical tide}}, \quad (9)$$

195 where

$$J_{\text{peak storm tide}} = \eta_{\text{storm tide}}(\mathbf{x}_0, t_{\text{peak storm tide}}), \quad (10)$$

196 and is computed with the model in full surge (tidally and meteorologically forced) mode, and

$$J_{\text{peak astronomical tide}} = \eta_{\text{tide only}}(\mathbf{x}_0, t_{\text{peak astronomical tide}}), \quad (11)$$

197 and is computed with the model in tide-only mode. The times $t_{\text{peak storm tide}}$ and $t_{\text{peak astronomical tide}}$ are
 198 determined from preliminary forward model runs. The sensitivity of the skew surge to each input m is then
 199 defined as

$$\frac{dJ_{\text{skew surge}}}{dm} = \frac{dJ_{\text{peak storm tide}}}{dm} - \frac{dJ_{\text{peak astronomical tide}}}{dm}. \quad (12)$$

200 A separate pair of forward and adjoint runs is required to evaluate the two terms on the right hand side
 201 of this expression; one with full (tide + meteorological) forcing, and one with tidal forcing only. However,
 202 when the model input m is taken as the wind stress, the second term on the right hand side is zero since,
 203 by definition, the tidally induced peak surface height is independent of the meteorological forcing.

204 Using the adjoint model, we evaluate equation (12) for three model input fields m , namely the bottom
 205 friction coefficient n , bathymetry h and wind stress $\boldsymbol{\tau}_s$. We note that the atmospheric pressure is also an
 206 important model input for surge modelling; however, for this event its overall contribution to the modelled
 207 storm tide is around 10%, and we therefore choose to focus on wind stress as the primary surge-generating
 208 input. The bottom friction coefficient and bathymetry are both scalar fields which are constant with respect
 209 to time, and the sensitivity pattern we compute with respect to these inputs is therefore only spatially
 210 varying. When computing the model sensitivities with respect to these inputs, we perform the adjoint model
 211 run over a period of approximately 10 days prior to the peak storm tide. This 10-day period was found to be
 212 sufficient, with longer periods having negligible effect on the computed sensitivities. Wind stress is a vector
 213 field which varies in both space and time, and the sensitivity of modelled skew surges with respect to wind
 214 stress is therefore also a spatially and temporally varying vector field. The wind stress sensitivity results
 215 presented here are computed from adjoint model runs spanning two days prior to the peak storm tide. The
 216 sensitivity of the modelled skew surge to the wind stress prior to this period was found to be small.

217 The spatial (and temporal) patterns of sensitivity to each model input reveal insights into the modelled
 218 system, but the sensitivities to different inputs cannot be directly compared, since they have incommensurable
 219 units. However, if we consider a perturbation in the input, Δm , and perform a convolution with the
 220 sensitivity, we can obtain an estimate of the resulting perturbation in the skew surge, ΔJ , via

$$\Delta J \approx \int \int \int \frac{dJ}{dm}(x, y; t) \cdot \Delta m \, dx \, dy \, dt. \quad (13)$$

221 This ΔJ can be directly compared for different inputs m , and we are thus able to compare the first-order
222 influence of each input parameter on the modelled skew surges, based on simple estimates for Δm . Note
223 that this is equivalent to performing a first-order Taylor expansion with respect to the input m , and that the
224 resulting ΔJ may be positive or negative for a given perturbation Δm . The use of this equation is valid as
225 long as the perturbations fall within the linear response regime of the model. While the inclusion of higher
226 order terms would extend the validity of a Taylor expansion approach beyond the range of linear response,
227 the calculation of higher order derivatives via adjoint methods is beyond the scope of this work.

228 For the purpose of uncertainty quantification within this paper, only spatially uniform Δm are considered.
229 This is an approximation to the true uncertainty, since real errors in the model inputs are likely to vary
230 spatially. However, in the absence of information about the spatial correlation of the input errors (the
231 estimation of which is beyond the scope of this work), ΔJ is taken as a simple estimate of the model output
232 uncertainty.

233 5. Results

234 5.1. Sensitivity to bottom friction coefficient

235 The fields of sensitivity to bottom friction coefficient for each target location are shown in figure 7.
236 The greatest sensitivity magnitudes are found within relatively small regions in the vicinity of each target
237 location. The sensitivity of the skew surge at North Shields exhibits the smallest sensitivity magnitudes, due
238 to its position on an exposed section of coastline; the propagation of the surge as a coastally trapped wave
239 is not strongly affected by local features, and the local value of the bottom friction coefficient therefore has
240 only a weak effect on the skew surge at the North Shields tide gauge. The sensitivity of the skew surge at
241 Immingham exhibits the greatest magnitudes, particularly in and around the Humber Estuary and the Wash.
242 The dynamics of the surge propagation around this region are complex, and the waters here are particularly
243 shallow; the $1/H$ proportionality in the wind stress and bottom stress terms of the governing equations (1)
244 therefore increases the model's sensitivity to bottom friction, as well as to bathymetry and wind stress, in
245 shallow waters. This high sensitivity to the friction coefficient in the region of the Humber Estuary and the
246 Wash is also evident for the skew surge at Lowestoft, suggesting that the interaction between the surge and
247 this region of coastline has a lasting effect on the surge as it travels further south.

248 Common to the sensitivity patterns for all target locations is the pattern in the far-field, i.e. in the north
249 of the domain. This is because any effect of the bottom friction on the surge in the north of the domain is
250 propagated with the surge as it travels south as a coastally trapped wave, and therefore has the same effect
251 on the skew surge at all target locations.

252 In order to estimate the total impact of an uncertain bottom friction coefficient on model outputs via
253 equation (13), we first estimate the uncertainty in the bottom friction coefficient (Δm in equation (13)).
254 Based on typical values for the Manning coefficient (Arcement and Schneider, 1989) for the types of sediment
255 found in the North Sea (Digimap, b), we assume an uncertainty in the Manning coefficient of $0.005 \text{ s m}^{-1/3}$.

256 Using equation (13) to convolve a uniform perturbation of $0.005 \text{ s m}^{-1/3}$ with the adjoint sensitivities shown
257 in figure 7, we obtain skew surge perturbations of -8.1 cm at North Shields, -17.3 cm at Immingham and
258 -16.1 cm at Lowestoft, with the minus signs indicating that increases in friction would induce reductions in
259 skew surge, due to the extraction of energy from the surge. The uncertainty in the mean skew surge along the
260 coastline section, estimated by the same method, is -19.9 cm; this is of similar magnitude to the estimated
261 uncertainties at Immingham (which is within the coastline section) and Lowestoft (just to the south of the
262 coastline section).

263 5.2. Sensitivity to bathymetry

264 The sensitivities of modelled skew surges to bathymetry are shown in figure 8. The observed spatial
265 patterns share similar features to those of the sensitivity to bottom friction coefficient of figure 7. We find
266 the greatest magnitudes of sensitivity within localised regions around each target location, and in particular
267 we find that these localised sensitivities share similar spatial patterns with those observed for bottom friction,
268 but with opposite signs. In the north of the domain, we again find that the observed patterns of bathymetry
269 sensitivity are similar for all target locations, due to the propagation of the surge as a coastally trapped wave
270 from north to south; any influence of the bathymetry on the surge in the north of the domain is propagated
271 south with the surge and impacts all subsequent observation locations.

272 To estimate the impact of this sensitivity on model outputs, we again start by estimating the uncer-
273 tainty in the bathymetry itself. For this, we compute the root mean square (RMS) difference between two
274 bathymetric datasets. We compare the GEBCO bathymetry dataset used within the model with data from
275 Digimap (Digimap, a), which is available at higher resolution than GEBCO, but does not cover the entire
276 model domain. In the region of our model domain in which both GEBCO and Digimap datasets are avail-
277 able, the RMS difference between the two is 2.7 m. Convolution of a uniform 2.7 m bathymetry perturbation
278 with the adjoint-computed bathymetry sensitivities via equation (13) produces perturbations of -2.3 cm, 6.7
279 cm and -4.8 cm in the skew surges at North Shields, Immingham and Lowestoft, respectively, and -3.7 cm
280 in the mean skew surge along the coastline section. The minus signs for North Shields, Lowestoft and the
281 coastline section indicate that an increase in bathymetry (i.e. an increase in water depth) induces a decrease
282 in the skew surge, with the opposite being the case at Immingham.

283 One feature common to the bathymetry sensitivity for Immingham, Lowestoft and the coastline section
284 is the region of positive sensitivity coinciding with Dogger Bank, to the north-east of the Humber Estuary
285 (see figure 1). The depth of this sand bank is around 20 m, with depths in excess of 60 m immediately
286 north of the bank. The positive sign of the bathymetry sensitivity in this region indicates that an increase
287 in bathymetry (i.e. the removal of the bank) would produce an increase in the skew surges at Immingham,
288 Lowestoft and the coastline section, and therefore that the bank protects the coastline to its south from the
289 surge.




Figure 7: Sensitivity of modelled skew surges to the bottom friction coefficient. Units: $\text{m s}^{-1} \text{m}^{1/3} \text{m}^{-2}$ (metres of surge, per unit Manning coefficient, per unit area). The relevant tide gauge locations are indicated by yellow circles. Top left: North Shields. Top right: Immingham. Bottom left: Lowestoft. Bottom right: mean along coastline section. The greatest sensitivity magnitudes are local to each location, and all locations exhibit similar patterns of sensitivity to bottom friction coefficient in the north of the domain.

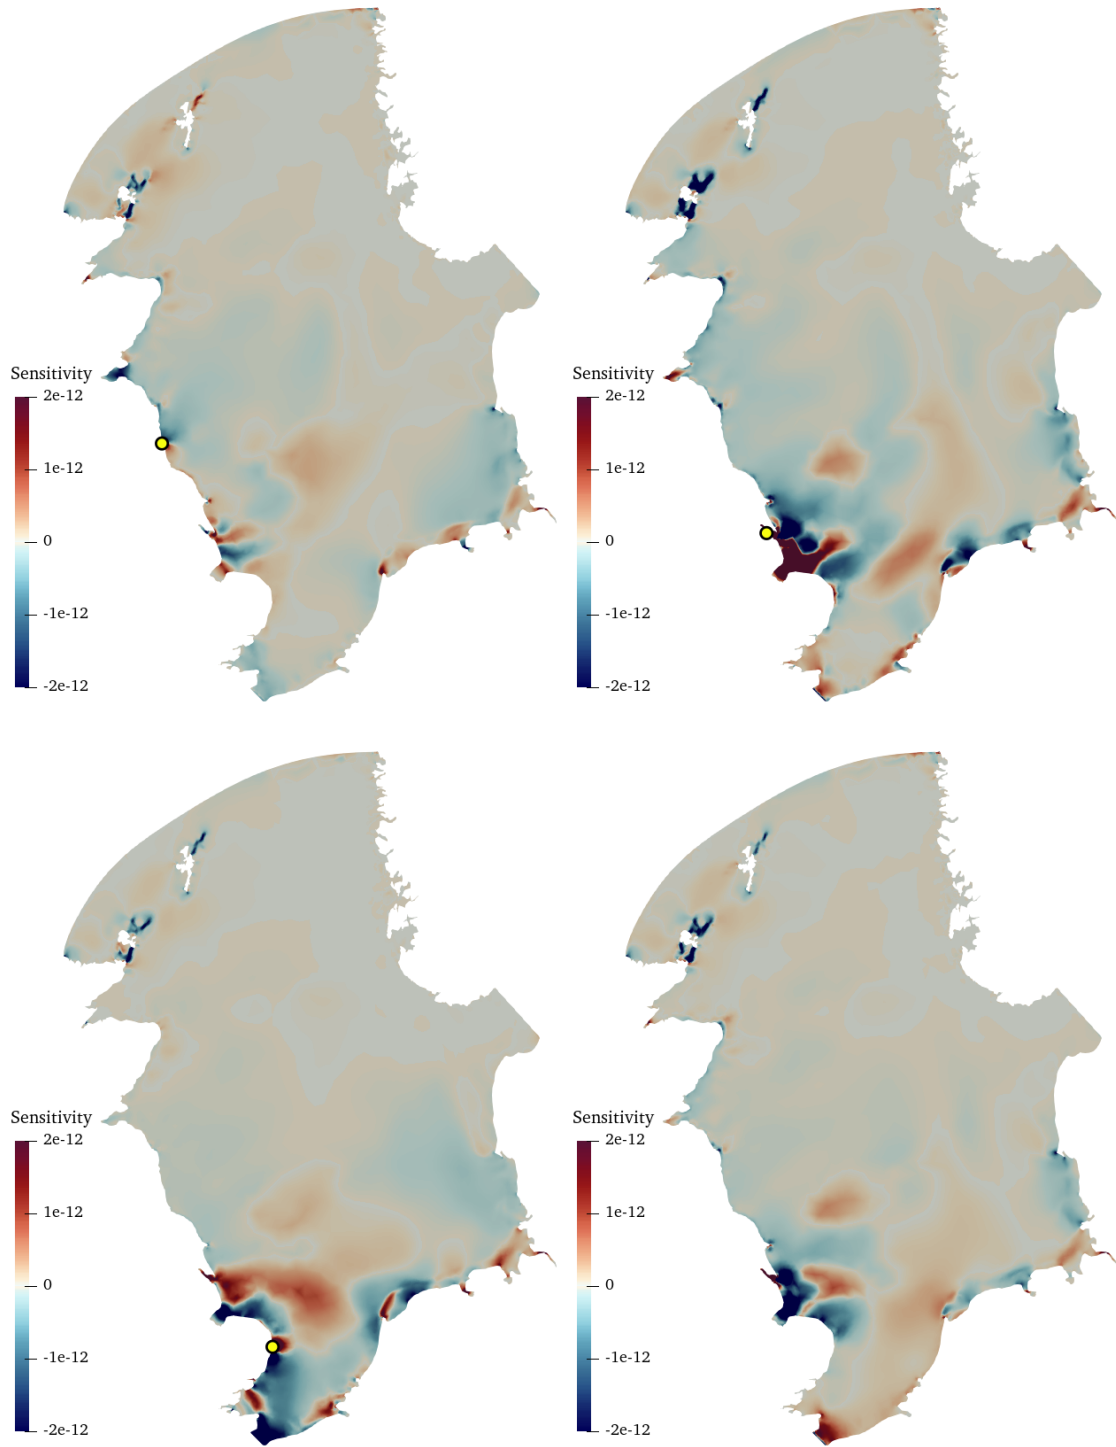


Figure 8: Sensitivity of modelled skew surges to the bathymetry. Units: $\text{m m}^{-1} \text{m}^{-2}$ (metres of surge, per metre of bathymetry, per unit area). The relevant tide gauge locations are indicated by yellow circles. Top left: North Shields. Top right: Immingham. Bottom left: Lowestoft. Bottom right: mean along coastline section. The greatest magnitudes are found in the vicinity of the target locations, and the patterns in the north of the domain are similar for all target locations.

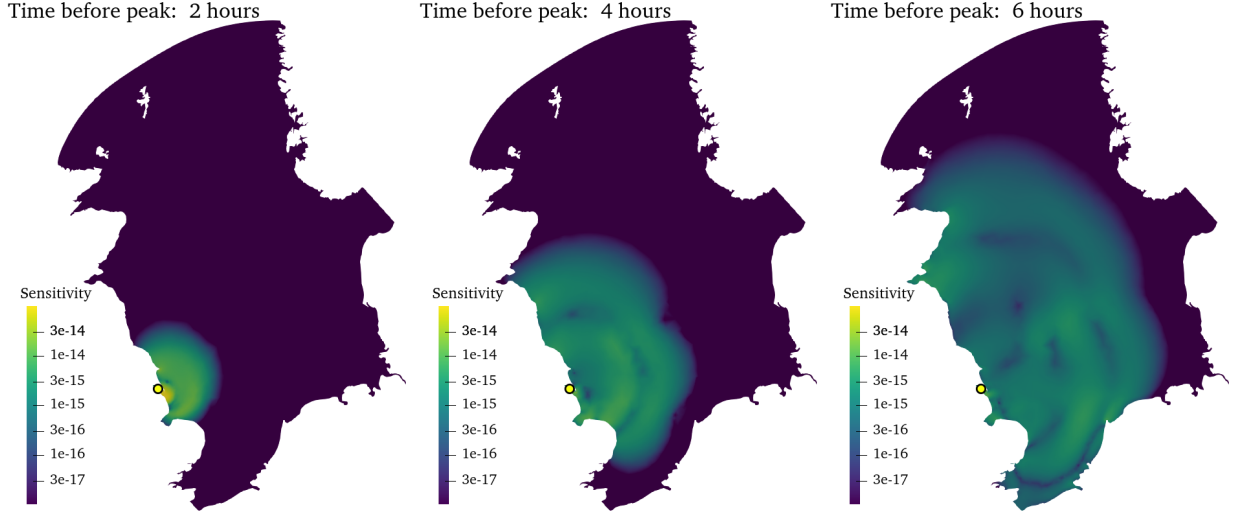


Figure 9: Snapshots of the magnitude of the sensitivity of the skew surge at Immingham to wind stress at various times prior to the peak storm tide, as labelled. Units: $\text{m Pa}^{-1} \text{m}^{-2} \text{s}^{-1}$ (metres of surge, per Pa wind stress, per unit area, per second). The yellow circle indicates the Immingham tide gauge location. The region of influence of the wind stress on the skew surge increases with lead time, due to the propagation of perturbations being limited by the shallow water wave speed.

290 5.3. Sensitivity to wind stress

291 Wind stress and atmospheric pressure are responsible for surge generation, and the sensitivity of a surge
 292 model to these inputs therefore has the potential to provide physical insight into the surge generation mech-
 293 anism. In an operational scenario, the meteorological inputs also carry high uncertainty, and understanding
 294 model sensitivity to these inputs is therefore essential to the interpretation of surge forecasts. Since the
 295 wind stress varies in space and time, so too do the sensitivities of modelled skew surges with respect to the
 296 wind stress. Considering a model output functional J corresponding to the peak storm tide elevation at a
 297 single location, the region of influence of the wind stress on J will expand as lead time increases, as shown
 298 in figure 9. This is due to the fact that the propagation of perturbations caused by wind stress is limited to
 299 the shallow water wave speed. For this reason, the sensitivity to wind stress can be considered as a shallow
 300 water wave propagating backwards in time, originating at the point at which the functional is defined. This
 301 has been explored in detail previously (Wilson et al., 2013), and can be further confirmed by an analytic
 302 approach (Warder et al., 2019).

303 In order to make progress comparing the wind stress sensitivities of skew surges at different locations, we
 304 can integrate the wind stress sensitivity field with respect to time to obtain an overall spatial pattern. These
 305 time-integrated sensitivities are shown in figure 10 for each target location. Similarly to the sensitivities to
 306 bottom friction coefficient and bathymetry, there are regions of high sensitivity magnitude in the vicinity of
 307 each target location, where local winds shortly before the peak storm tide occurs have a significant effect
 308 on the value of the peak sea surface height (and hence skew surge). All four target locations exhibit similar
 309 patterns of sensitivity to wind stress in the north of the domain, but differ more in the south, because any

310 perturbations induced by wind stress in the north of the domain affect the coastally trapped wave which
311 then impacts all target locations as it travels south.

312 The magnitudes of wind stress sensitivity are generally greater in the west of the domain. This is due
313 to the southerly propagation of the surge along the western coastal boundary of the model domain (the
314 east coast of the UK); winds in the east of the domain therefore have relatively little effect on the surge
315 impacting the UK locations considered within this study. In particular, sensitivity magnitudes over the
316 Norwegian Trench are very small. This is likely to be due to the very large depths in this region, and the
317 $1/H$ proportionality in the wind stress term in the governing equations (1).

318 It is not possible to make a generally applicable estimate of the uncertainty associated with wind stress,
319 since in a forecast scenario this depends strongly on the forecast lead time, and the nature of the surge
320 event. To make a simple comparison between the overall wind stress contribution to uncertainty for each
321 target location, we take uniform (in both space and time) wind stress perturbations of 0.1 Pa in each of the
322 x - and y -directions, and convolve these with the adjoint-derived sensitivities via equation (13). The output
323 perturbations calculated for the 0.1 Pa wind stress perturbations in the positive x -direction are 0.5, -1.2,
324 3.0 and 0.3 cm for North Shields, Immingham, Lowestoft and the coastline section, respectively. The corre-
325 sponding output perturbations for wind stress perturbations in the positive y -direction are -5.5, -8.0, -10.8
326 and -9.1 cm, respectively. Firstly, these results show that modelled skew surges can be significantly increased
327 by wind stress perturbations in the negative y -direction (north to south), while wind stress perturbations in
328 the x -direction (east-west) have a smaller effect which is more variable across the target locations. Secondly,
329 the overall sensitivities to wind stress perturbations show an increasing trend for gauges further south, due
330 to the southward propagation of the surge and the corresponding accumulation of influence of wind stress.

331 6. Discussion

332 6.1. Comparison of uncertainties

333 In section 5 we used the adjoint model to explore the spatial patterns of storm surge model sensitivity to its
334 uncertain inputs. In the cases of bottom friction coefficient and bathymetry, we have estimated uncertainties
335 in each model input and, through convolution with the model sensitivity, estimated the resulting uncertainties
336 in the model outputs, namely the skew surges at selected coastal target locations. In contrast to the raw
337 sensitivities, these estimated output uncertainties can be directly compared. A summary of these estimated
338 uncertainties is shown in table 1. We make three key observations:

339 (i) Estimated uncertainties due to bottom friction are of greater magnitude than those due to bathymetry.

340 However, it should be noted that we have made very simple estimates of input uncertainties, and for a
341 well-calibrated model these uncertainties would likely be significantly reduced. This is particularly the
342 case for the bottom friction, and these results highlight the importance of achieving a tight constraint
343 on the bottom friction coefficient through model calibration methods.

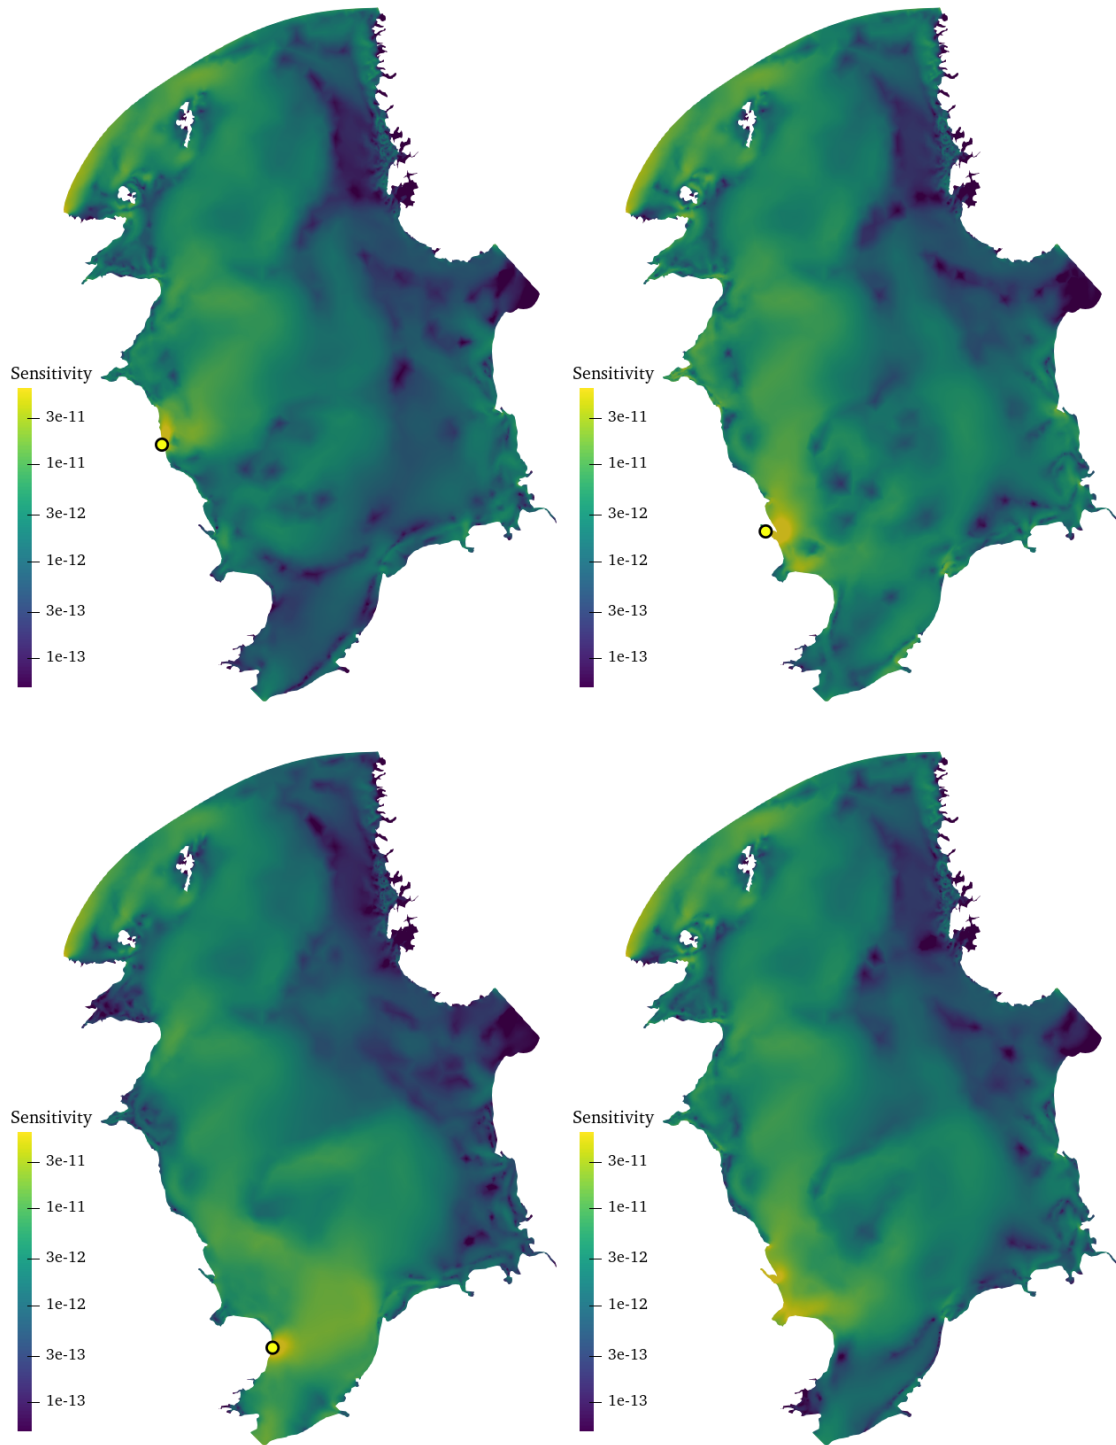


Figure 10: The magnitude of the time-integrated sensitivity of modelled skew surges to wind stress. Units: $\text{m Pa}^{-1} \text{m}^{-2}$ (metres of surge, per Pa wind stress, per unit area). The relevant tide gauge locations are indicated by yellow circles. Top left: North Shields. Top right: Immingham. Bottom left: Lowestoft. Bottom right: coastline section. The greatest magnitudes are local to each target location, and there is a similar pattern in the north of the domain for all target locations.

	North Shields	Immingham	Lowestoft	Coastline section
Bottom friction coefficient ($\pm 0.005 \text{ s m}^{-1/3}$)	$\mp 8.1 \text{ cm}$	$\mp 17.3 \text{ cm}$	$\mp 16.1 \text{ cm}$	$\mp 19.9 \text{ cm}$
Bathymetry ($\pm 2.7 \text{ m}$)	$\mp 2.3 \text{ cm}$	$\pm 6.7 \text{ cm}$	$\mp 4.8 \text{ cm}$	$\mp 3.7 \text{ cm}$

Table 1: Summary of estimated skew surge uncertainties due to bottom friction coefficient and bathymetry, calculated from adjoint-derived sensitivities and estimated input uncertainties. In almost all cases, the response of the skew surge to positive perturbations in the inputs is a decrease in the skew surge, as indicated by the \mp signs in the uncertainties, i.e. deeper water or increased friction results in decreased skew surges. The effect of bathymetry at Immingham is the exception, where a positive bathymetry perturbation (deeper water) results in increased skew surge.

- 344 (ii) The uncertainty contributed by the bottom friction is of smaller magnitude for the northernmost target
345 location (North Shields) than for the locations further south, which all exhibit similar magnitudes. This
346 pattern is explained by the accumulation of uncertainty over the propagation path of the surge along
347 the east coast of the UK; bottom friction acts to remove energy from the surge, and this effect is
348 therefore cumulative along the path of the surge from north to south.
- 349 (iii) In contrast, the overall contribution of uncertain bathymetry exhibits a more variable pattern across
350 the domain, suggesting that the effect of the bathymetry on the skew surge arises through a variety of
351 mechanisms. The similarity (with opposite signs) between the localised spatial patterns of sensitivity to
352 bottom friction coefficient and bathymetry (figures 7 and 8) suggests that a proportion of the sensitivity
353 to bathymetry in these regions arises from the bottom friction term of the governing equations, which
354 is inversely proportional to the water depth. However, given the contrasting patterns of estimated
355 uncertainty due to each input summarised in table 1, it is clear that the sensitivity to bathymetry is
356 more complex, and must also derive significant contributions from the other terms of the governing
357 equations (1) in which the bathymetry h appears, i.e. the wind stress and surface elevation advection
358 terms.

359 Regarding the overall sensitivity to wind stress, our results indicate that modelled skew surges exhibit
360 positive sensitivity to wind stress perturbations in the negative y -direction, i.e. that wind stress perturbations
361 aligned with the southerly propagation of the surge act to increase the peak storm tide. Since the uncertainty
362 in the wind stress depends strongly on the forecast lead time, a direct comparison between uncertainty due
363 to bottom friction, bathymetry and wind stress is not possible. However, we know from ensemble forecasts
364 for this event that the uncertainty due to meteorological inputs was on the order of 1 m at a forecast lead
365 time of 24 hours; this is far greater than the uncertainties due to bottom friction and bathymetry estimated
366 here. The quantitative results of this study are therefore consistent with the perceived limitations of the
367 operational model at the time, namely that storm surge forecast model performance is limited by the accuracy
368 of the meteorological forecast providing the wind stress (and atmospheric pressure).

369 *6.2. Implications and limitations*

370 The results of an adjoint sensitivity analysis as performed within this study are highly relevant at the
371 interface between models and observations. The spatial pattern of sensitivity to bottom friction coefficient
372 could, for example, be used to inform the intelligent application of a spatially varying bottom friction
373 coefficient, for the purposes of more sophisticated model calibration. For example, a choice of length scale of
374 variation in bottom friction coefficient could be made based on the spatial variability of the model sensitivity,
375 since variations on smaller length scales would not be constrained by observations. Similarly, the relatively
376 localised impact of uncertain bathymetry shown here suggests that bathymetric surveys, particularly in
377 regions prone to morphological change, could be valuable in reducing uncertainty in storm surge forecasts.
378 This bathymetric sensitivity also suggests that the impact of imposing a minimum water depth for model
379 stability purposes, as we have done within this study, should be carefully considered. Finally, the observed
380 patterns of sensitivity to wind stress could be used to inform efforts to enhance meteorological models, by
381 identifying regions in which uncertainty in wind stress has the greatest impact on overall surge uncertainty.

382 In addition to assisting in analysing surge model performance, the adjoint-based sensitivity analysis
383 performed within this work is capable of providing physical insight into surge generation and propagation.
384 The skew surges at Immingham, Lowestoft and the coastline section all show a positive gradient with respect
385 to the bathymetry over Dogger Bank, to the north-east of the Humber Estuary; this is visible in figure 8.
386 This reveals the protective effect of this bank for the south-east coast of the UK, against this storm surge
387 event. Similarly, the sensitivity to wind stress of figure 10 shows very low sensitivity over the Norwegian
388 Trench, due to the deep water in this region. Features such as these are simple to interpret within the physics
389 contained in the governing equations. However, quantifying the impact of these features on the generation
390 and propagation of the surge is non-trivial, but is achieved at relatively low computational cost by the adjoint
391 techniques employed here.

392 The sensitivity analysis approach we have taken here consists of computing gradients of model outputs
393 with respect to model inputs. This facilitates a linearisation of the model with respect to the inputs con-
394 sidered, i.e. the use of a Taylor expansion as a substitute for the full forward model, via equation (13).
395 This expansion is only valid for sufficiently small perturbations of the model inputs, but could be used, for
396 example, to estimate an arbitrarily large ensemble of model outputs at the computational cost of only one
397 forward and one adjoint model run (since the cost of evaluating the Taylor expansion is negligible compared
398 to running the full numerical model). Since the adjoint model used here requires approximately 2.4 times
399 the computation time of the forward model, this constitutes a highly efficient approach. This is of particular
400 interest for uncertain wind stress, where operational uncertainty quantification is typically carried out using
401 ensemble methods. However, the viability of the adjoint-based approach as a substitute for an ensemble
402 method is limited by two key factors. Firstly, the range of perturbations in the ensemble may exceed the
403 linear response regime of the model, and secondly, the adjoint model must be computed separately for every
404 model output of interest. An ensemble forecast or hazard assessment over a large spatial scale is therefore not

405 feasible using adjoint methods alone. However, as we have shown here, an adjoint-based sensitivity analysis
406 can provide information complementary to ensemble methods; for a given event, the adjoint in conjunction
407 with an ensemble method could provide a more complete analysis of the potential inundation consequences
408 for flood risk assessment purposes than ensemble methods alone.

409 We note a number of avenues for further work. In order to gain estimates of model uncertainties due to
410 each input, we have propagated uniform perturbations through the model, via the adjoint-derived sensitivi-
411 ties. The use of spatially varying input perturbations would require an estimate of the spatial correlation of
412 errors in the model inputs, which was beyond the scope of this work but may be considered in future. We
413 also note that the model output perturbations could have been computed from forward model runs alone.
414 However, the use of the Taylor expansion approach via the adjoint-derived sensitivities is a demonstration
415 of the efficiency of the adjoint method for propagating input perturbations through the numerical model at
416 low computational cost (once the adjoint model has been run), and is central to the efficient propagation of
417 large ensembles through the model, as described above.

418 We further note a number of modelling choices we have made within this study, whose influence could be
419 investigated in future work. Firstly, we have not included waves within our numerical model, even though
420 wave effects may have contributed 40 cm to the surge for this event (Staneva et al., 2017). In particular, we
421 have used a constant Charnock parameter to capture the atmosphere-ocean coupling, whereas this parameter
422 is thought to depend on wave age (Drennan et al., 2005, Brown and Wolf, 2009). However, our modelling
423 assumptions are consistent with the CS3X model used operationally at the time of the case study event,
424 and our model performance (based on the RMSEs of modelled surge residuals) is comparable with CS3X.
425 Furthermore, the model sensitivity to the wind stress is independent of the wind stress parameterisation itself
426 (i.e. we are not computing sensitivity with respect to the wind velocity), and we therefore assume that our
427 results are not significantly impacted by our choice to neglect waves. Similarly, we have not included near-
428 shore effects such as wave setup, which may have influenced the observed storm tide. We leave the inclusion
429 of these effects, and the study of their subsequent impact on model sensitivities, to future work. Secondly,
430 we have used a uniform value for the friction coefficient within the model, although the spatial variation of
431 this parameter can have a significant impact on modelled surges (as highlighted within this work). However,
432 additional forward model runs revealed that, for a perturbation in the Manning coefficient of $0.001 \text{ s m}^{-1/3}$,
433 nonlinear contributions to the resulting storm tide perturbation were around 2%. We therefore assume
434 that nonlinear effects are sufficiently small that the computed sensitivities to the friction coefficient do not
435 depend strongly on its input value. Finally, the use of a minimum water depth, to avoid the need to include
436 wetting and drying within the model, may inhibit the predictive capability of the model, and is likely to have
437 influenced the model calibration. Due to the model resolution and bathymetry dataset used, the application
438 of a 10 m minimum depth only impacts upon a small fraction of the UK east coast, although parts of the
439 Humber Estuary and the Wash were affected, thus potentially impacting the model results at Immingham
440 and the coastline section, and possibly Lowestoft. The results of this work, which show high sensitivity

441 to bathymetry in localised regions around each target location, suggest that the bathymetry modification
442 in shallow regions may be important, and therefore that accurate surge modelling requires small values of
443 imposed minimum depths, or ideally the inclusion of wetting and drying, even if inundation modelling is
444 not the focus. The use of a minimum depth within this study may have led to an underestimation of model
445 sensitivities to all uncertain inputs in these localised shallow regions, but an investigation into this influence
446 is beyond the scope of this study.

447 7. Conclusions

448 In this work, we have applied adjoint methods to perform sensitivity analysis and uncertainty quan-
449 tification for a storm surge model, in particular comparing the sensitivity of the modelled skew surge, at
450 different locations across the domain, to three different model inputs, namely the bottom friction coefficient,
451 bathymetry and wind stress. Based on the results of this work, conclusions can be drawn based on both the
452 underlying sensitivity patterns revealed, and also the resulting estimates of model uncertainty due to each
453 of the model inputs.

454 The underlying patterns of skew surge sensitivity to all model inputs considered exhibit high spatial
455 variability, with high sensitivity magnitudes in localised regions around each target location. However, we
456 also find that the sensitivity to model inputs in the north of the domain is similar for all target locations; i.e.
457 perturbations in bottom friction, bathymetry or wind stress in the north of the domain have a similar impact
458 on all target locations. This is consistent with the storm surge propagating south as a coastally trapped wave
459 along the east coast of the UK, since any effect of the model inputs on the surge in the north of the domain
460 will travel south with the wave and impact all locations in its path. The spatial variability of sensitivity
461 to each input has potentially broad implications, such as the application of a spatially varying bottom
462 friction coefficient, the commissioning of new bathymetric surveys in regions where high sensitivity aligns
463 with high bathymetry uncertainty, or to provide feedback informing improvements of the meteorological
464 models providing the wind stress and atmospheric pressure forcing for surge models.

465 Physical insight can also be gained from the patterns of surge sensitivity. For example, we see in the
466 sensitivity to bathymetry that locations on the UK coast towards the south of the domain are protected
467 from the surge by Dogger Bank, a large sand bank around 200 km off the UK coast. We also find that
468 sensitivity to wind stress is particularly low over the Norwegian Trench, due to the very deep water. These
469 are good examples of how adjoint methods can be used to gain physical insight, and form a valuable tool for
470 analysing the impact of a storm surge event.

471 Using the adjoint-derived sensitivities to estimate the uncertainty in a skew surge model prediction
472 due to typical uncertainty in each input, we find that an uncertainty of $0.005 \text{ s m}^{-1/3}$ in the Manning
473 coefficient produces uncertainty of around 15 cm in the modelled skew surge, highlighting the importance
474 of model calibration in constraining this uncertainty. Similarly, we estimate that an uncertainty of 2.7 m in
475 the bathymetry produces uncertainty of around 5 cm in the modelled skew surge. The contribution from

476 uncertain meteorological inputs can be on the order of 1 m in an operational forecast scenario, far exceeding
477 the uncertainty due to bottom friction or bathymetry, and ensemble methods remain the most practical
478 approach to uncertainty quantification in a forecast scenario. However, we have shown here how an adjoint-
479 based sensitivity analysis provides complementary information to an ensemble approach, providing detailed
480 spatial and temporal information about how input uncertainty is mapped onto outputs.

481 Acknowledgements

482 This work was funded by the EPSRC Centre for Doctoral Training in Fluid Dynamics across Scales (Grant
483 EP/L016230/1). MDP would additionally like to acknowledge EPSRC support under Grant EP/R029423/1.
484 We thank Jane Williams of the National Oceanography Centre for her contribution of hindcast data. We also
485 acknowledge the Research Computing Service at Imperial College London for access to computing resources.
486 This study uses data from the National Tidal and Sea Level Facility, provided by the British Oceanographic
487 Data Centre and funded by the Environment Agency.

488 References

- 489 G. J. Arcement and V. R. Schneider. Guide for Selecting Manning’s Roughness Coefficients for Natural
490 Channels and Flood Plains. Technical report, 1989.
- 491 ASC. Managing climate risks to well-being and the economy. Adaptation Sub-Committee. Committee on
492 Climate Change. Progress report 2014., 2014.
- 493 A. Avdis, A. S. Candy, J. Hill, S. C. Kramer, and M. D. Piggott. Efficient unstructured mesh generation for
494 marine renewable energy applications. *Renewable Energy*, 116:842–856, 2018. ISSN 0960-1481. doi: <https://doi.org/10.1016/j.renene.2017.09.058>. URL <http://www.sciencedirect.com/science/article/pii/S0960148117309205>.
- 497 J. M. Brown and J. Wolf. Coupled wave and surge modelling for the eastern Irish Sea and implica-
498 tions for model wind-stress. *Continental Shelf Research*, 29(10):1329–1342, may 2009. ISSN 0278-
499 4343. doi: 10.1016/J.CSR.2009.03.004. URL <https://www.sciencedirect.com/science/article/pii/S0278434309001010>.
- 501 R. Canizares, A. W. Heemink, and H. J. Vested. Application of advanced data assimilation methods for the
502 initialisation of storm surge models. *Journal of Hydraulic Research*, 36(4):655–674, 1998. doi: 10.1080/
503 00221689809498614. URL <https://doi.org/10.1080/00221689809498614>.
- 504 H. Chen, A. Cao, J. Zhang, C. Miao, and X. Lv. Estimation of spatially varying open boundary
505 conditions for a numerical internal tidal model with adjoint method. *Mathematics and Computers
506 in Simulation*, 97:14–38, mar 2014. ISSN 0378-4754. doi: 10.1016/J.MATCOM.2013.08.005. URL
507 <https://www.sciencedirect.com/science/article/pii/S0378475413002000>.

508 R. Comblen, J. Lambrechts, J.-F. Remacle, and V. Legat. Practical evaluation of five partly discontinuous
509 finite element pairs for the non-conservative shallow water equations. *International Journal for Numerical*
510 *Methods in Fluids*, 63(6):701–724, 2010.

511 Digimap. Marine Themes Digital Elevation Model 6 Arc Second [ASC geospatial data], Scale 1:250000, Tiles:
512 2062000000, 2060010060, 2060010040, 2060010020, 2060000000, 2058010080, 2058010060, 2058010040,
513 2058010020, 2058000000, 2056010080, 2056010060, 2056010040, 2056010020, 2056000020, 2056000000,
514 2054010080, 2054010060, 2054010040, 2054010020, 2054000020,2054000000, 2052010080, 2052010060,
515 2052010040, 2052010020, 2052000020, 2052000000, 2050010080, 2050010060, 2050010040, 2050010020,
516 2050000020, 2050000000, 2048010080, 2048010060, 2048010040, Updated: 25 October 2013, Ocean-
517 Wise, Using: EDINA Marine Digimap Service, <https://digimap.edina.ac.uk>, Downloaded: 2019-10-25
518 12:26:52.575, a.

519 Digimap. DiGSBS250K [SHAPE geospatial data], Scale 1:250000, Tiles: GB, Updated: 6 September 2011,
520 BGS, Using: EDINA Geology Digimap Service, <https://digimap.edina.ac.uk>, Downloaded: 2019-12-05
521 15:24:22.171, b.

522 W. M. Drennan, P. K. Taylor, and M. J. Yelland. Parameterizing the sea surface roughness. *Journal of*
523 *physical oceanography*, 35(5):835–848, 2005.

524 G. D. Egbert and S. Y. Erofeeva. Efficient inverse modeling of barotropic ocean tides. *Journal of Atmospheric*
525 *and Oceanic Technology*, 19(2):183–204, 2002. ISSN 07390572. doi: 10.1175/1520-0426(2002)019(0183:
526 EIMOBO)2.0.CO;2.

527 P. E. Farrell, D. A. Ham, S. W. Funke, and M. E. Rognes. Automated derivation of the adjoint of high-level
528 transient finite element programs. *SIAM Journal on Scientific Computing*, 35(4):C369–C393, 2013.

529 J. Flowerdew, K. Horsburgh, C. Wilson, and K. Mylne. Development and evaluation of an ensemble forecast-
530 ing system for coastal storm surges. *Quarterly Journal of the Royal Meteorological Society*, 136:1444–1456,
531 2010. ISSN 00359009. doi: 10.1002/qj.648.

532 J. Flowerdew, K. Mylne, C. Jones, and H. Titley. Extending the forecast range of the uk storm surge
533 ensemble. *Quarterly Journal of the Royal Meteorological Society*, 139(670):184–197, 2013.

534 S. W. Funke. *The automation of PDE-constrained optimisation and its applications*. PhD thesis, Imperial
535 College London, 2012.

536 C. Geuzaine and J. F. Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-
537 processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009.
538 ISSN 00295981. doi: 10.1002/nme.2579.

539 Z. Goss, D. Coles, and M. Piggott. Identifying economically viable tidal sites within the alderney race
540 through optimization of levelized cost of energy. *Philosophical Transactions of the Royal Society A*, 378
541 (2178):20190500, 2020.

542 I. D. Haigh, M. P. Wadey, T. Wahl, O. Ozsoy, R. J. Nicholls, J. M. Brown, K. Horsburgh, and B. Gouldby.
543 Spatial and temporal analysis of extreme sea level and storm surge events around the coastline of the UK.
544 *Scientific Data*, 3:1–14, 2016. doi: doi:10.1038/sdata.2016.107.

545 M. C. G. Hall, D. G. Cacuci, and M. E. Schlesinger. Sensitivity Analysis of a Radiative-Convective Model
546 by the Adjoint Method, 1982. ISSN 0022-4928.

547 S. Hallegatte, C. Green, R. J. Nicholls, and J. Corfee-Morlot. Future flood losses in major coastal cities.
548 *Nature Climate Change*, 3(9):802–806, 2013. ISSN 1758678X. doi: 10.1038/nclimate1979.

549 A. Heemink, E. Mouthaan, M. Roest, E. Vollebregt, K. Robaczewska, and M. Verlaan. Inverse 3D shallow wa-
550 ter flow modelling of the continental shelf. *Continental Shelf Research*, 22(3):465–484, feb 2002. ISSN 0278-
551 4343. doi: 10.1016/S0278-4343(01)00071-1. URL [https://www.sciencedirect.com/science/article/
552 pii/S0278434301000711](https://www.sciencedirect.com/science/article/pii/S0278434301000711).

553 T. Kärnä, S. C. Kramer, L. Mitchell, D. A. Ham, M. D. Piggott, and A. M. Baptista. Thetis coastal ocean
554 model: Discontinuous Galerkin discretization for the three-dimensional hydrostatic equations. *Geoscientific Model Development*, 11(11):4359–4382, 2018. ISSN 19919603. doi: 10.5194/gmd-11-4359-2018.

556 S. Kramer, T. Kärnä, J. Hill, and S. W. Funke. stephankramer/uptide: First release of uptide v1.0, 2020.
557 <http://doi.org/10.5281/zenodo.3909652>.

558 R. W. Lardner, A. H. Al-Rabeh, and N. Gunay. Optimal Estimation of Parameters for a Two-Dimensional
559 Hydrodynamical Model of the Arabian Gulf. *Journal of Geophysical Research*, 98(C10):18229–18242, 1993.

560 Y. Li, S. Peng, J. Yan, and L. Xie. On improving storm surge forecasting using an adjoint optimal technique.
561 *Ocean Modelling*, 72:185–197, dec 2013. ISSN 1463-5003. doi: 10.1016/J.OCEMOD.2013.08.009. URL
562 <https://www.sciencedirect.com/science/article/pii/S1463500313001625>.

563 M. Losch and P. Heimbach. Adjoint sensitivity of an ocean general circulation model to bottom topography.
564 *Journal of Physical Oceanography*, 37(2):377–393, 2007.

565 X. Lu and J. Zhang. Numerical study on spatially varying bottom friction coefficient of a 2D tidal model with
566 adjoint method. *Continental Shelf Research*, 26(16):1905–1923, oct 2006. ISSN 0278-4343. doi: 10.1016/J.
567 CSR.2006.06.007. URL <https://www.sciencedirect.com/science/article/pii/S027843430600210X>.

568 S. Massmann. Sensitivities of an adjoint, unstructured mesh, tidal model on the European Continental
569 Shelf. *Ocean Dynamics*, 60(6):1463–1477, dec 2010. ISSN 16167341. doi: 10.1007/s10236-010-0347-6.
570 URL <https://doi.org/10.1007/s10236-010-0347-6>.

571 S. K. Mitusch, S. W. Funke, and J. S. Dokken. dolfin-adjoint 2018.1: automated adjoints for FEniCS and
572 Firedrake. *Journal of Open Source Software*, 4(38), 2019.

573 W. Nowak. Using algorithmic differentiation for uncertainty analysis. In *22nd Telemac & Mascaret User*
574 *Club*, pages 52–57, 2015. doi: 10.5281/zenodo.165522.

575 C. Pain, M. Piggott, A. Goddard, F. Fang, G. Gorman, D. Marshall, M. Eaton, P. Power, and C. De Oliveira.
576 Three-dimensional unstructured mesh ocean modelling. *Ocean Modelling*, 10(1-2):5–33, 2005.

577 F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. McRae, G. T. Bercea, G. R. Markall,
578 and P. H. Kelly. Firedrake: Automating the finite element method by composing abstractions. *ACM*
579 *Transactions on Mathematical Software*, 43(3), 2016. ISSN 15577295. doi: 10.1145/2998441.

580 J. Staneva, V. Alari, Ø. Breivik, J.-R. Bidlot, and K. Mogensen. Effects of wave-induced forcing on a
581 circulation model of the north sea. *Ocean Dynamics*, 67(1):81–101, 2017.

582 T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M.
583 (eds.). *IPCC, 2013: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I*
584 *to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University
585 Press, Cambridge, United Kingdom and New York, NY, USA, 2013.

586 A. Verdy, M. R. Mazloff, B. D. Cornuelle, and S. Y. Kim. Wind-driven sea level variability on the california
587 coast: An adjoint sensitivity analysis. *Journal of physical oceanography*, 44(1):297–318, 2014.

588 C. Villaret, R. Kopmann, D. Wyncoll, J. Riehm, U. Merkel, and U. Naumann. First-order uncertainty anal-
589 ysis using Algorithmic Differentiation of morphodynamic models. *Computers & Geosciences*, 90:144–151,
590 may 2016. ISSN 0098-3004. doi: 10.1016/J.CAGEO.2015.10.012. URL <https://www.sciencedirect.com/science/article/pii/S0098300415300777>.

592 C. V. Vouriot, A. Angeloudis, S. C. Kramer, and M. D. Piggott. Fate of large-scale vortices in idealized
593 tidal lagoons. *Environmental Fluid Mechanics*, 19(2):329–348, apr 2019. ISSN 15731510. doi: 10.1007/
594 s10652-018-9626-4. URL <https://doi.org/10.1007/s10652-018-9626-4>.

595 M. P. Wadey, I. D. Haigh, R. J. Nicholls, J. M. Brown, K. Horsburgh, B. Carroll, S. L. Gallop, T. Mason, and
596 E. Bradshaw. A comparison of the 31 January–1 February 1953 and 5–6 December 2013 coastal flood events
597 around the UK. *Frontiers in Marine Science*, 2:84, 2015. ISSN 2296-7745. doi: 10.3389/fmars.2015.00084.
598 URL <https://www.frontiersin.org/article/10.3389/fmars.2015.00084>.

599 S. C. Warder, K. J. Horsburgh, and M. D. Piggott. Understanding the contribution of uncertain wind stress
600 to storm surge predictions. In *4th IMA International Conference on Flood Risk*, Swansea, 2019.

601 J. Williams and R. Flather. Interfacing the operational storm surge model to a new mesoscale atmospheric
602 mode. 2000.

- 603 J. Williams, K. J. Horsburgh, J. A. Williams, and R. N. Proctor. Tide and skew surge independence:
604 New insights for flood risk. *Geophysical Research Letters*, 43(12):6410–6417, 2016. ISSN 19448007. doi:
605 10.1002/2016GL069522.
- 606 C. Wilson, K. J. Horsburgh, J. Williams, J. Flowerdew, and L. Zanna. Tide-surge adjoint modeling: A
607 new technique to understand forecast uncertainty. *Journal of Geophysical Research: Oceans*, 118(10):
608 5092–5108, 2013. ISSN 21699291. doi: 10.1002/jgrc.20364.
- 609 J. Zhang, X. Lu, P. Wang, and Y. P. Wang. Study on linear and nonlinear bottom friction parameterizations
610 for regional tidal models using data assimilation. *Continental Shelf Research*, 31(6):555–573, apr 2011.
611 ISSN 0278-4343. doi: 10.1016/J.CSR.2010.12.011. URL [https://www.sciencedirect.com/science/
612 article/pii/S0278434310003857](https://www.sciencedirect.com/science/article/pii/S0278434310003857).