Earth ArXiv

01

## RESEARCH ARTICLE

# Finding the potential height of tropical cyclone storm surges in a changing climate using Bayesian optimization

Simon D.A. Thomas[1,2] [iD]*, Dani C. Jones[2,3] [iD], Talea Mayo[4] [iD], John R. Taylor[1], Henry B. Moss[1,5], David R. Munday[2] [iD], Ivan D. Haigh[6] [iD] and Devaraj Gopinathan[7] [iD]

[1]Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, Cambridgeshire, UK
[2]British Antarctic Survey, United Kingdom Research and Innovation, Cambridge, Cambridgeshire, UK
[3]Cooperative Institute for Great Lakes Research, University of Michigan, Ann Arbor, Michigan, USA
[4]Department of Mathematics, Emory University, Atlanta, Georgia, USA
[5]School of Mathematics, Lancaster University, Lancaster, Lancashire, UK
[6]Department of Oceanography, University of Southampton, Southampton, Hampshire, UK
[7]Advanced Research Computing Centre, University College London, London, Greater London, UK
*Corresponding author. Email: sdat2@cam.ac.uk

## Abstract

We introduce a new framework for systematically exploring the largest storm surge heights that a tropical cyclone in a given climate can create. We calculate the tropical cyclone potential intensity and the potential size from climate model projections and find that both these limits increase in response to climate change. We then use Bayesian optimization with a barotropic ocean circulation model to find the maximum height that the surge can reach given these limits. The key methodological advances of this paper are (i) calculation of the recently proposed potential size of a tropical cyclone, now and under climate change (ii) using Bayesian optimization to find the largest storm surge given those constraints, (iii) using this information to constrain the return level curve. This paper uses key theoretical improvements in our understanding of tropical cyclones to understand implications for changing storm surge risk. We have chosen the US coastline and the area around New Orleans as our case study area, but this method is generalizable and could in principle be applied to any coastline.

### Impact Statement

Our methodology provides a way of assessing the impact of climate change on storm surges. We calculate previously proposed upper bounds on the intensity and size of a tropical cyclone given the climate, and then use machine learning to provide an efficient method of calculating the storm surge. By focusing on the worst case storm surge and its relation to the climatic conditions, we can improve predictions of damaging long return period storm surges. Our novel approach can easily be transferred to other coastlines around the world that are influenced by tropical cyclones.

## 1. Introduction

Whilst state-of-the-art climate models are invaluable research tools for many questions, they struggle to explicitly resolve tropical cyclones. As climate models must parameterize processes that require spatial scales smaller than the lateral grid box spacing (typically 1° in CMIP6 as in CESM2, Danabasoglu *et al.* (2020)), these small-scale processes (e.g. the eye-wall, rainbands and warm ocean currents) risk being misrepresented (Camargo & Wing 2016)). Furthermore, the creation of tropical cyclones (cyclogenesis) relies on seeding events such as African easterly waves (AEWs), and a bias in the simulation of these leads to a bias in cyclogenesis (Camargo & Wing 2016). Similarly, there are well-known biases in tropical cyclone seasonality (Sainsbury *et al.* 2022; Peng & Guo 2024; Shan *et al.* 2023), frequency (Sainsbury *et al.* 2022) intensification (Roberts *et al.* 2020), and intensity (Roberts *et al.* 2020) within modern climate models. When storm surge models are then forced with tropical cyclones from these models, even with attempts to bias correct them, these biases can be propagated into estimates of the storm surge hazard and resultant risk (e.g. Sobel *et al.* (2023)).

Tropical cyclones are also dependent on broad-scale fields such as surface temperature. Along with a finite amplitude genesis event (such as an AEW) (Emanuel 1991), tropical cyclones are generated due to a thermodynamic disequilibrium between the sea surface and the tropopause (Emanuel 1986; Emanuel 2003; Emanuel 2006). The potential intensity (PI) was introduced by Emanuel (1986), by imagining a Carnot cycle running between these reservoirs, and has been long-accepted as a reasonable measure for the upper bound in tropical cyclone azimuthal speeds (Rousseau-Rizzi *et al.* 2021). The potential intensity is expected to increase with climate change as the disequilibrium between reservoirs is increased by the enhanced greenhouse effect, and there is observational evidence that the upper bound of tropical cyclone intensity has increased over recent years (Wehner & Kossin 2024). Studies have previously used the potential intensity to investigate subsequent limits of resulting storm surge heights (e.g. Mori *et al.* (2022)). The size of the tropical cyclone also significantly affects the size of a storm surge, and Wang *et al.* (2022a) proposed a new measure of potential size (PS), though it has not yet been used with climate model output or to constrain downstream consequences. Given these two limits (i.e. of potential intensity and potential size), the height of storm surges generated by tropical cyclones should be limited, too.

In order to relate these limits of potential size and intensity to a corresponding limit on storm surge height, we can use a storm surge model ADCIRC (Luettich Jr & Westerink 1991) running in a Bayesian Optimization loop. We assume that both PI and PS are achieved at the same time, taking a value from the grid-point nearest that point along the coast at that time, and then as in Jia & Taflanidis (2013) we vary the tropical cyclone trajectory and speed to find the highest storm surge height that can be created at that point. This helps to incorporate the complex interplay between the bathymetry and storm surge trajectory in a flexible and efficient way, and Ide *et al.* (2024) showed that Bayesian optimization can effectively find the most impactful TC trajectory for a storm surge. Mori *et al.* (2022) used the calculated potential intensity together with coastal response functions (Irish *et al.* 2009) instead of a full storm surge model to show that the potential storm surge height increased, with a study area of bays in East Asia, but although using response functions is likely more efficient, this may not fully account for the effects of complex coastal geometry, and they were not able to include the potential size change in tropical cyclones from Wang *et al.* (2022a). On the other hand, Lin & Emanuel (2015) ran the ADCIRC model with very large catalogues of realistic tropical cyclones to estimate the long return period "Grey Swan" tropical cyclone storm surges for numerous points around the world. This is likely much more computationally intensive, but allows for more complex TC trajectories.

Understanding the worst-case storm surge, which we choose to call the potential height, is useful for several reasons. The design requirements for critical infrastructure and emergency response and planning efforts rely on estimates of low probability and worst-case scenarios. Some structures are designed to withstand very high return period events (e.g. coastal nuclear power plants Schwerdt *et*

***Table 1.*** *Abbreviations used during this paper*

| Abbreviation | Expansion |
|---|---|
| PI | Tropical cyclone potential intensity (Emanuel 1986) |
| PS | Tropical cyclone potential size (Wang *et al.* 2022a) |
| CLE15 | Chavas *et al.* (2015)'s TC wind profile |
| TC | Tropical cyclone |
| W22 | Wang *et al.* (2022a) Carnot engine |
| GP | Gaussian process |
| DAF | Data acquisition function |
| MES | Max-value entropy search (Wang & Jegelka 2017) |
| EVT | Extreme value theory |
| GEV | Generalized extreme value distribution |
| CMIP6 | Coupled Model Intercomparison Project Phase 6 |
| ADCIRC | ADvanced CIRCulation model |

*al.* (1979) and ONR (2014), which are designed to withstand a one in a 10,000 year storm surge), and having an estimate of the potential height may be an efficient way of informing the design requirements.

The framework we introduce here can be extended to other coastlines around the world that are influenced by tropical cyclones. Our work is novel in three ways:

1. Calculating tropical cyclone potential size using climate data, both in the present and future (which builds upon Wang *et al.* (2022a) and Wang *et al.* (2023)).
2. Using Bayesian optimization to find the worst possible storm surge from the potential intensity and potential size limits.
3. Using this worst-case estimate and extreme value theory to demonstrate how they can reduce uncertainties in the return period curve.

We first discuss our proposed framework (Section 2) and explain how the potential intensity (PI) and potential size (PS) thermodynamic limits are calculated (Section 2.2). We then describe the climate data used to calculate these limits, and the methods used to process them (Section 2.1). We explain how the ADCIRC model is used to calculate the storm surge (Section 2.3). In Section 2.4 we describe how Bayesian optimization is used to find the worst-case storm surge. Finally, we explain how extreme value theory is used to find the return period curve (Section 2.5). This method is shown as a flowchart in Figure 1. We present our results in Section 3. We first show how we can optimize for one point along the coast, and then demonstrate how we can extend our finding to many points along the coast to understand how the maxima and the characteristics of the maxima vary spatially and over time. We then discuss the utility and limitations of this approach (Section 4). We conclude our study and summarise our key results in Section 5. To aid readability, abbreviations have been listed in Table 1.

## 2. Methodology

### 2.1. Climate data

The potential intensity and potential size limits can be calculated based on the output of climate models, which will allow us to explore how these limits might evolve in the future. From the current state-of-the-art collection of climate models, CMIP6, we can use the historical experiments for plausible realisations of the past climate from 1850-2014, and the SSP-585 scenario, the most pessimistic scenario of those developed for 2014-2100. While SSP-585 may not be the most likely outcome given recent progress in renewable energy (Huard *et al.* 2022), it provides a large warming signal difference that should make it easier as a first test of our techniques. We have chosen to use the CMIP6 models included in the
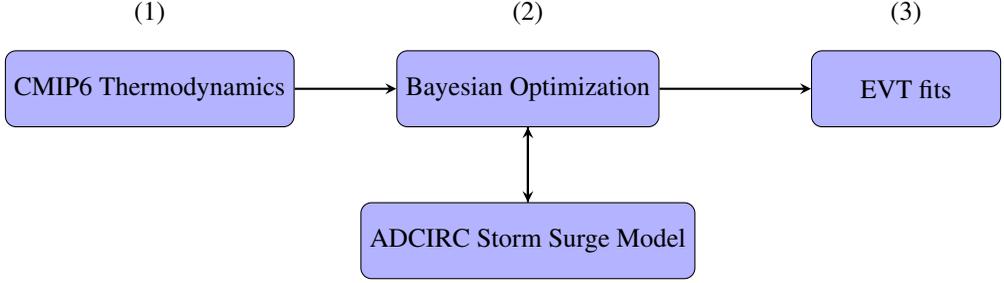
**Figure 1.** *Flowchart of the methodology. The climate data (Section 2.1) is used to calculate the potential intensity and potential size limits (CMIP6 thermodyanmics, Section 2.2). These are then used to constrain the idealized tropical cyclone input to the ADCIRC model (Section 2.3). The ADCIRC model is then used to calculate the storm surge by repetitively driving it with idealized tropical cyclones, whose track parameters are chosen within a Bayesian Optimisation loop (Section 2.4). The maximum storm surge limit is then used as a constraint to fit the extreme value theory distributions (Section 2.5)..*

Pangeo data catalogue's Google Cloud Store (Pangeo 2022) because this data is easily available. We have initially chosen to focus on an ensemble member of the CESM2 climate model (Danabasoglu *et al.* 2020) but the analysis could be simply extended to any of the other models in the catalogue. The data from the Pangeo catalogue is preprocessed using the xMIP preprocessing package (Busecke *et al.* 2023), which ensures the ensemble members are all converted to SI units and compatible label names. The data is then regridded to $\frac{1}{2}^{\circ}$ using CDO (Schulzweida 2023).Initially we have not implemented a method of bias correction either for the mean state or trend (as in Mori *et al.* (2022)), but we would like to explore the effects of bias correction in the future on both measures.

### 2.2. Potential intensity and potential size

We use the tropical cyclone potential intensity first proposed in Emanuel (1986), in its most common formulation as a measure of the maximum azimuthal windspeed attainable (as reviewed in Rousseau-Rizzi *et al.* (2021)). This is calculated following Bister & Emanuel (2002) as

$$\left(V_p\right)^2 = \frac{T_s}{T_0} \frac{C_k}{C_D} \left(\text{CAPE}^* - \text{CAPE}_{\text{env}}\right)|_{RMW}, \tag{PI}$$

where $V_p$ is the potential intensity at the gradient wind level ($V_{\text{reduc}} = 1$ in Gilford (2021)), CAPE$^*$ is the convective available potential energy of saturated air lifted from the sea surface to the outflow level, CAPE$_{\text{env}}$ is the convective available potential energy of the environment, $C_K$ is the enthalpy exchange parameter, $C_D$ is the momentum exchange parameter, $T_s$ is the surface temperature, and $T_o$ is the outflow temperature, and it it evaluated at the radius of maixmum winds, RMW. The tcPyPI provides a user-friendly package for calculating the potential intensity (Gilford 2021), and we assume $\frac{C_K}{C_D} = 0.9$ which is standard.

   Despite the acceptance of the potential intensity of a tropical cyclone, there has been less success in determining how large a tropical cyclone can grow. The potential size metric was suggested by Wang *et al.* (2022a) based on assuming both an improved Carnot cycle running between the sea surface and the tropopause and the azimuthally symmetric wind and pressure dynamical profile of Chavas *et al.* (2015). This is reasonable because as the storm becomes larger, the pressure drop modeled by the Carnot cycle is reduced, and the pressure drop assumed by the dynamical profile with cyclogeostrophic balance increases (following Chavas *et al.* (2017)). There will be one radius of outer winds where both models predict the same pressure drop at a particular storm outer radius (Wang *et al.* 2022a). This solution can be found using a method such as the bisection method, finding the point at which the two modeled radii

intersect. We assume that the maximum windspeed from the Chavas *et al.* (2015) profile is the potential intensity $V_p$. We then assume that the maximum windspeed for the Wang *et al.* (2022a) W22 Carnot engine, $V_{\text{max W22}}$, is given by the potential intensity multiplied by the supergradient factor $\gamma_{\text{sg}} = 1.2$ of Wang *et al.* (2022a). We are therefore assuming that the cyclone is both at its maximum potential intensity $V_p$ and maximum potential size $r_a$ at the same time.

This potential size $r_a$ is justified by assuming that if a tropical cyclone has reached the maximum windspeed under the potential intensity limit, there might still be additional energy (work) that could be extracted by expanding the size of the tropical cyclone. The point at which the maximum windspeed of the tropical cyclone can be maximally expanded, whilst satisfying both the dynamical limits from assuming the Chavas *et al.* (2015) radial profile and assuming the Wang *et al.* (2022a)'s updated tropical cyclone Carnot engine is called the potential size $r_a$. In their paper this is calculated as the outer radius of vanishing winds ($r_a$), but the Chavas *et al.* (2015) profile with a given velocity of maximum winds $V_p$ has a single radius of maximum winds, $r_{\text{max}}$, that this would correspond to assuming that the other environmental conditions are held constant. Further details of the potential intensity and size calculation procedure are discussed in Appendix 6.

Potential size depends on potential intensity, and both of these quantities depend on the broad-scale climate variables of the atmospheric profile and surface variables at a point. Figures 2(a) & (c) show the potential intensity and size calculated on a processed monthly average from a single CMIP6 ensemble member (CESM2-r4i1p1f1) for August 2015. As shown in Figure 2(a), potential intensity is higher in areas with higher sea surface temperature (a spatial correlation coefficient of $\rho = 0.97$ between SST and potential intensity, and a linear gradient fit of $m = 11 \pm 1$ m s$^{-1}$ °C$^{-1}$). Whereas, Figure 2(c) shows that geographically potential size is dominated by the north-south contrast (a spatial correlation coefficient of $\rho = -0.99$ between latitude and potential size ($m = -88 \pm 1$ km °N$^{-1}$)), consistent with the $\frac{V_p}{f}$ scaling in Wang *et al.* (2022a) where $f$ is the Coriolis parameter and $V_p$ is the potential intensity ($\rho = 0.92, m = (1.48 \pm 0.03) \times 10^3$ m rad$^{-1}$). Therefore, as might be expected, much larger and more intense tropical cyclones are possible in the warmer waters further south.

When we look at a point near New Orleans point near New Orleans (29.25°N, -90.25°E) where we have plotted the metrics calculated from monthly average data from CESM2-r4i1p1f1 (black) and CESM2-r10i1p1f1 (purple) for August the SSP-585 (2015-2100) scenarios, the two metrics vary over time in quite different ways (Figure 2b&d). Potential intensity (Figure 2(b)) shows much higher variability relative to its magnitude than potential size (Figure 2d). For r4i1p1f1, over the SSP-585 scenario from 2015-2100 the increase of potential size ($\rho = 0.89, m = 2.3 \pm 0.1$ km yr$^{-1}$) is more visible and significant than potential intensity ($\rho = 0.57, m = (6.6 \pm 1.0) \times 10^{-2}$ m s$^{-1}$ yr$^{-1}$), where $\rho$ is the correlation coefficient against time, and $m$ is the gradient of a linear fit over the period, and we calculate the standard deviation of this parameter based on the fit. Therefore, both of these metrics appear to increase as a response to climate change, but potential size seems to show a more significant effect, which is an interesting property not highlighted in Wang *et al.* (2022a). There is a much stronger correlation between the timeseries during SSP-585 at this point for SST and potential size ($\rho = 0.93, m = 57 \pm 3$ km °C$^{-1}$) compared to SST and potential intensity ($\rho = 0.70, m = 1.9 \pm 0.2$ m s$^{-1}$ °C$^{-1}$).

### 2.3. ADCIRC model driven by idealized tropical cyclone

We use the ADCIRC model in the barotropic 2D depth-integrated mode (Luettich *et al.* 1992; Westerink *et al.* 1994), which is a state-of-the-art storm surge model solved on an unstructured mesh with triangular elements. ADCIRC solves the generalized wave continuity equation using a Jacobi preconditioned iterative solver (Westerink *et al.* 1994). Many different meshes are available online to resolve the coastline in a variety of different levels of detail. We used the EC95c mesh (see e.g. Dietrich *et al.* (2013)), which resolves the East and Gulf of Mexico US coastlines using 58,369 triangular elements with 31,435 triangular vertices (nodes). This resolution captures some small-scale structures of the US such as the barrier islands on the coast. The model lacks some geographic features such as Lake Pontchartrain to the north of New Orleans, and does not allow the water to flow onto the land. We input the atmospheric
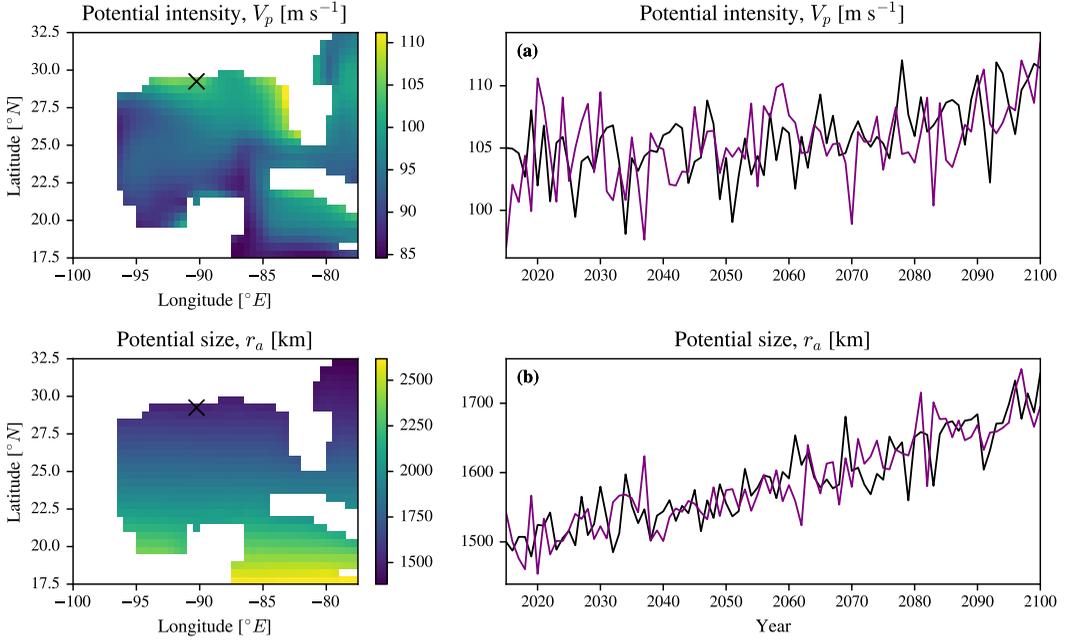
**Figure 2.** *Tropical cyclone potential intensity and potential size calculated on CESM2-r10i1p1f1 ensemble member for the historical period and the SSP-585 (high-emission) scenario. The top left and top right panels are of potential intensity, and the bottom left and bottom right panels are for potential size. The left panels show the variables calculated for a single snapshot based on monthly average properties for August 2015 (from near the start of SSP-585). The right panels show the same indices calculated for based on the August monthly average conditions for years between 2015 and 2100 for the grid point near New Orleans highlighted on the left-hand panel for r4i1p1f1 (black) and r10i1p1f1 (purple) .*

forcing fields using the netCDF format, with a large stationary grid over the whole domain at a resolution of $\frac{1}{8}^{\circ}$, and a moving grid centered on the tropical cyclone center at a resolution of $\frac{1}{80}^{\circ}$. Both grids are regular, orientated to be parallel to lines of constant latitude and longitude. The ADCIRC setup does not include tides, and has a timestep of 5 seconds. Further details of the storm surge model mesh and its settings are described in Appendix 7. An example snapshot of the model being forced by the CLE15 wind profile for August 2015 is shown in Figure 3.

### *2.4. Bayesian Optimization and Surrogate Modelling*

In Bayesian optimization (Garnett 2023), we seek to find the global optimum of a function $f(\vec{x}^*)$, where $f : \vec{x} \in \mathbb{R}^N \rightarrow z \in \mathbb{R}$. We seek to do this with a minimum number of samples of $\vec{x}$, $N_s$, so as to reduce the computational cost of finding $f(\vec{x}^*)$, especially when $f$ is computationally expensive. To make repeated evaluations of $f$ amenable, we replace $f$ with its inexpensive statistical surrogate $\hat{f}$ (also known as a meta-model, or emulator). To reduce the total number of samples, as well as to find samples that are most informative of $f$, we define a data acquisition function (DAF), $\alpha_t(\vec{x})$. The DAF maximises mutual information to find the most informative next sample or samples for $f$ (or $\hat{f}$) after it has been fitted on all of the previous samples.

Bayesian optimization has been widely used in applications where each additional sample is computationally expensive to utilize, e.g. in changing the hyper– and meta– parameters of a deep learning
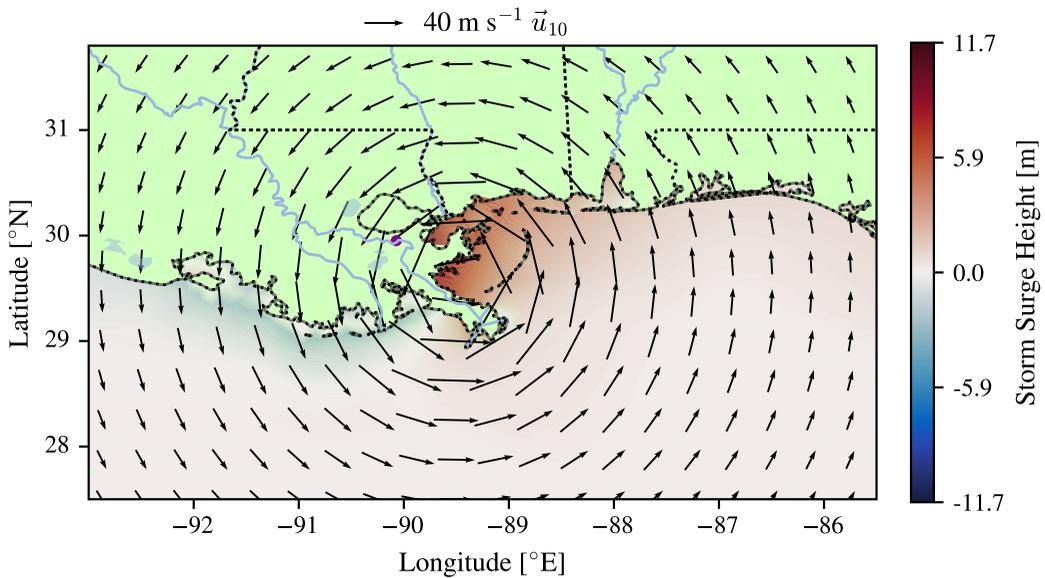
**Figure 3.** *A snapshot from the 31,435 node ADCIRC model driven by an idealized "worst case" hurricane for August 2015, with the CLE15 windfield at the potential size and potential intensity for that month. The quivers show the 10m wind velocity vectors with size proportional to magnitude, and the colormap shows the storm surge height. The purple point is the centre of the city of New Orleans. The land is shaded green, with the rivers and lakes shaded blue. The rivers/lakes/land are not part of the model domain at this stage, but they could be included in ADCIRC. The model domain includes the whole of the Gulf of Mexico and the East Coast of the United States. The dashed lines mark the borders of US states, including their coastline, but the model's coastline is less detailed than these dashed lines .*

model (Garnett 2023), where retraining a model may take days of computational time. It is also a good approach to replicate simulations from numerical models of physical systems, which can be very expensive and often depend on many parameters that are hard to tune (see e.g. Khatamsaz *et al.* (2023)).

In our case, rather than tuning the parameters of the numerical storm-surge model, we aim to change the characteristics of the idealized tropical cyclone used to force it. Defining the limits of the tropical cyclone characteristics $\vec{x}$, we can sequentially run ADCIRC with different sets of tropical cyclone characteristics and produce values of the resulting maximum storm surge heights at a point $z$. To be consistent with the derivation of the potential size, we use the azimuthally symmetric wind profile of Chavas *et al.* (2015) to define the tropical cyclone that forces the ADCIRC model in terms of the tropical cyclone characteristics. We choose to set the wind profile's maximum velocity, $V_p$, and outer radius, $r_a$, to the potential intensity and potential size limit calculated for the respective year. These could also be varied but to reduce the degrees of freedom we assume that the largest surge will occur during the most intense and largest storm. In addition, we assume that the tropical cyclone travels on a line of constant bearing and speed, which has three degrees of freedom: the bearing of the trajectory $\chi$, displacement east and west of the trajectory $c$, and the translation speed $V_t$ (similar parameters to those of Hashemi *et al.* (2016)). We limit our modelling to the simulation of surge, and omit tides to eliminate the variability due to the tidal cycle.

We develop a surrogate of the ADCIRC model that efficiently maps the input degrees of freedom $\vec{x} = (\chi, c, V_t)$ to the output, i.e. the maximum storm surge height at a point in the domain, $z$. Specifically, we

use Gaussian process (GP) emulators within our Bayesian optimization framework because GPs work well with small amounts of data in a low-dimensional parameter space and provide predictions of the mean and associated Gaussian uncertainty (Williams & Rasmussen 2006). A GP is fully characterized by its mean, $\mu(\vec{x})$, and covariance function, $k(\vec{x}_i, \vec{x}_j)$, where the prior mean $\mu(\vec{x})$ is normally taken to be 0 and the covariance function is chosen to embody the prior understanding of how the data behave. We choose to use the Matérn (5/2) kernel,

$$k(\vec{x}_i, \vec{x}_j) = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left( \frac{\sqrt{2\nu}}{l} d(\vec{x}_i, \vec{x}_j) \right)^{\nu} K_{\nu} \left( \frac{\sqrt{2\nu}}{l} d(\vec{x}_i, \vec{x}_j) \right), \tag{1}$$

where $\nu = \frac{5}{2}$, $\Gamma(\cdot)$ is the gamma function, $K_{\nu}(\cdot)$ is a modified bessel function, $l$ is the length scale, and $d$ is Euclidian distance between $\vec{x}_i$ and $\vec{x}_j$. We make this choice because it was found to be the most performant in a low-resolution version of ADCIRC in terms of reducing the RMSE and negative log-likelihood for a test set. This is consistent with other studies (e.g. Gopinathan *et al.* (2021)).

To initially fit the surrogate model, $\hat{f}$, we first sample $L$=25 points using Latin hypercube search (LHS), which is a space filling design that places points away from each other (McKay *et al.* 1979). After this we use the max-value entropy search acquisition function $\alpha_t$ (MES, Wang & Jegelka (2017)), which is particularly good at finding the global optimum of a function that is expensive to evaluate, to sequentially select an additional $B = 25$ Bayesian optimization points to evaluate, refitting the surrogate model $\hat{f}$ after each additional sample. This is shown as a flowchart in Figure 4.

The number of initial points $L$ and additional points $B$ to select is chosen arbitrarily, although we show in Figure 5 that using the MES DAF leads to substantially better optima for one particular point than continuing to sample points using LHS. In strategy (A) we just use a latin hypercube search (LHS) to fill the space with 50 samples ($L = 50$, $B = 0$). In strategy (B) we first use 25 LHS samples and then 25 samples using the maximum value entropy search (MES) data acquisition function ($L = 25$, $B = 25$). As defined, the two strategies are equivalent up until the 25th sample. To show the difference between the two strategies for optimization, in Figure 5 we plot the approximate simple regret, which we define as the difference between the current maximum observed storm surge, $\max\left(\vec{z}^i_{1,\cdots,s}\right)$, up to that sample index $s$ in the trial index $i$ and the maximum of all trials (1 to $n$) over all samples (the stand in for the global maximum), $\max\left(\max\left(\vec{z}^1\right), \cdots \max\left(\vec{z}^n\right)\right)$, so that

$$\text{Approximate Simple Regret}_{si} = \max\left(\max\left(\vec{z}^1\right), \cdots \max\left(\vec{z}^n\right)\right) - \max\left(\vec{z}^i_{1,\cdots,s}\right). \tag{2}$$

This approximate simple regret is a measure of how much the maximum of the trials is behind the approximate global maximum, and is used in place of the simple regret (see e.g. Wang *et al.* (2022a)). As expected, the two strategies appear indistinguishable with the trials of the two experiments overlapping before the 25th sample (dashed line), because both select the same 25 LHS points. By 25 additional points (50 points total) then all trials of using Bayesian optimization with the MES DAF in strategy (B) outperform continuing with additional LHS space filling points in strategy (A). Therefore, this suggests that we can be confident that strategy (B) used for later experiments is superior to a simple LHS design, at least for this point on the coast, for finding a higher optima with equal computational resources. This does not show that strategy (B) is the optimal strategy, and it is possible that a different ratio of LHS to DAF samples, a different DAF (e.g. Expected Improvement as in Ide *et al.* (2024)), and an improved GP kernel (see e.g. Tazi *et al.* (2023)) could all improve the performance of the optimization strategy, where the best setting for each could depend on the point on the coast.

An example of the Bayesian Optimization process (continuing to follow strategy (B), as in the rest of this paper) is shown in Figure 6, where, for clarity, we vary just the tropical cyclone track's bearing and displacement east and west of New Orleans. We use ADCIRC to compute the maximum storm surge height at the point in the mesh closest to New Orleans, and define this as the storm surge height. In
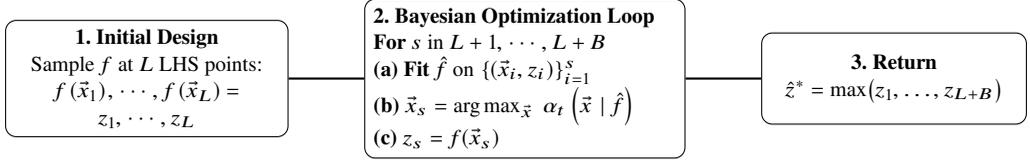
**Figure 4.** *Bayesian optimization flow chart for an experiment with L initial LHS samples and B Bayesian optimization samples. In our case $f$ is the ADCIRC model wrapped to take an idealized TC input with a particular velocity and pressure profile $V(r)$, $p(r)$, and variable TC track parameters $\vec{x} = (c, \chi, V_t)$, returning the maximum SSH $z$ over the simulation at a single observation node. $\hat{f}$ is the GP emulator and $\alpha_t$ is the MES DAF. The highest sample during the experiment is returned.*
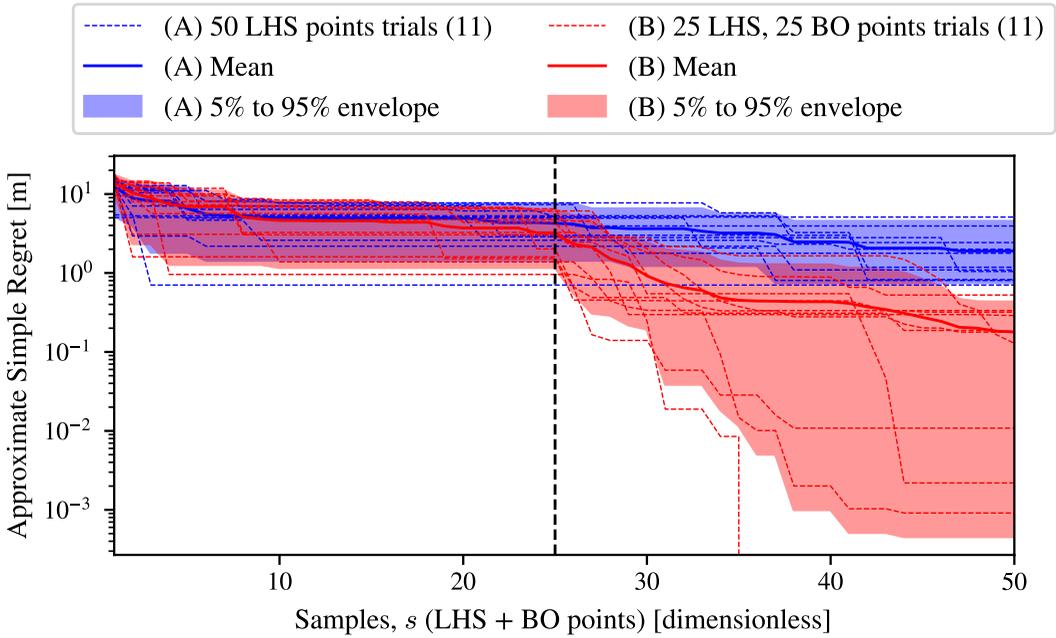


**Figure 5.** *Plotting the approximate simple regret to compare the strategy of (A) 50 initial LHS samples for the point in the mesh closest to New Orleans, and (B) 25 initial LHS samples followed by 25 MES DAF samples. For each strategy 11 trials with different random seeds are used. The dashed lines show the current maximum in each trial at that sample s. The solid lines show the mean of the maximum of the trials at that sample s. The shaded regions show the 5% to 95% estimate of the trials at that sample s. Up until 25 samples (black dashed vertical line) the two strategies are equivalent.*

all of the panels, we plot the initial 25 points/ADCIRC simulations selected through Latin Hypercube sampling as blue crosses (x), and show the three additional points selected by Bayesian optimization as green plusses (+). Figure 6(a) shows the GP emulator's mean, $\mu_{\hat{f}}(\vec{x})$, and Figure 6(b) shows the emulator's standard deviation $\sigma_{\hat{f}}(\vec{x})$, after it has been fitted on all of the points sampled so far. In Figure 6(b) the standard deviation of the emulator is much lower around the points which have already been sampled. We can see that the emulator expects the highest maximum storm surge heights to be reached either at positive bearing and positive track displacement or negative bearing and negative track displacement, forming a band of higher values (Figure 6(a)). The two green points that have already
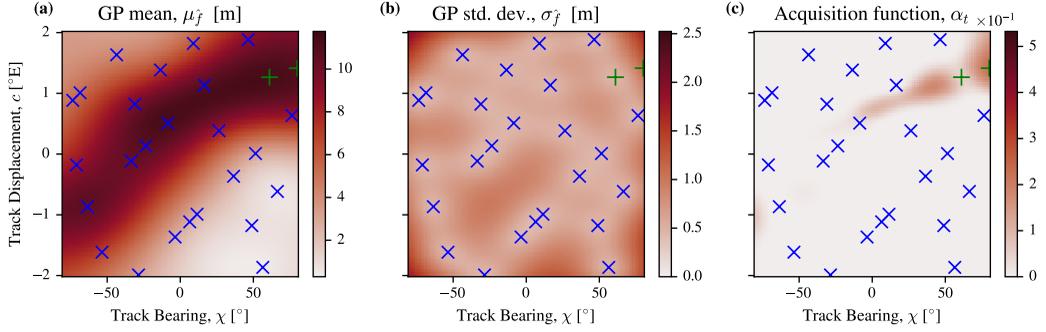
**Figure 6.** *An example of active learning in two-dimensional space of hurricane trajectory, using the Matérn 5/2 kernel, where the units for each graph are the size of the storm surge in meters at a node point to the north of New Orleans. The blue crosses in each panel are the 25 initial training points selected by Latin hypercube design, and the green pluses are the first three additional samples selected. Panel (a) shows the emulator's prediction of the maximum storm surge height at each point in the domain, and panel (b) shows the emulator's uncertainty at each point in the domain. Panel (c) shows the data acquisition function which in this case is the expected improvement. Subsequent points are chosen as the maximum of this function.*

been sampled are at positively bearing and positive track displacement, suggesting this end of the track-bearing ridge has the highest emulator values $\hat{f}$. The MES DAF in Figure 6(c) is highest near the three additional points that have already been selected, so that the model will continue to explore this area of the parameter space for the next sample.

### 2.5. Fitting a Generalized Extreme Value Distribution

In the final part of the framework, we explore the use of the information provided by Bayesian optimization in constraining the fit of a statistical extreme value model. To demonstrate the value of knowing the upper bound, we conduct a set of idealized statistical simulations.

   If we consider the case where we have block-maxima observations (the maximum storm surge height recorded each year), then we would expect that we can model the dataset $\vec{z}$ of $N_s$ observations as samples from a Generalized Extreme Value (GEV) distribution (Coles *et al.* 2001). If we assume the existence of a maximum $z^*$ that the distribution should reach corresponding to the potential height, we can assume the distribution is of the Weibull class. For the case of the GEV distribution the probability density function is given by

$$\text{GEV}(z; \alpha, \beta, \gamma) = \frac{1}{\beta}\left(1 + \gamma\left(\frac{z-\alpha}{\beta}\right)\right)^{-1-\frac{1}{\gamma}}\exp\left(-\left(1 + \gamma\left(\frac{z-\alpha}{\beta}\right)\right)^{-\frac{1}{\gamma}}\right), \tag{3}$$

where $z$ is the maximum height of the storm surge above the coast for a particular year, $\alpha$ is the position parameter, $\beta$ is the scale parameter, and $\gamma$ is the shape parameter, which should be $\gamma < 0$ for the Weibull class distribution. We can then find the upper bound $z^* = -\frac{\beta}{\gamma} + \alpha$.

   In the case where we exactly know the upper bound ahead of time (case I) this leaves us with two parameters to fit, $\beta$ and $\gamma$. If we do not know the upper bound ahead of time (case II) and have to fit it, we have three parameters to fit $\alpha$, $\beta$, and $\gamma$. In both cases we can minimize the negative log-likelihood of the dataset $\vec{z}$ with respect to the parameters using tensorflow with the Adam optimizer (Kingma 2014) for 1,000 optimization steps with a learning rate of 0.01, $\beta_1 = 0.9$, $\beta_2 = 0.99$, and $\epsilon = 10^{-6}$ which were the default parameters.

To assess the recovery of the original distribution with these two cases (I and II), we take $N_s$ samples from a GEV distribution, and choose the parameters $\alpha = 2$m, $\beta = 1$m, and $\gamma = -0.2$, and therefore $z^* = 7$m. To estimate the uncertainty in the model parameters we resample the dataset $N_r = 600$ times for each setting and refit both cases each time. This was chosen empirically, as we found $N_r = 600$ to be large enough to reliably estimate confidence intervals. In order to focus on metrics of interest for risk professionals we focus on the 1 in 100 year (0.01% exceedance probability per year) and 1 in 500 year (0.002% exceedance probability per year) return values calculated from the fitted distributions. For each case using these $N_r = 600$ estimates, we calculate the mean and the 5% and 95% for both return values.

Because of calculation error, and simplifications made in each part of the Potential Height framework, we will not know the upper bound exactly. We can make case I more realistic by adding some Gaussian noise to the true upper bound $z^*$ so that the 'calculated' upper bound is $\hat{z}^* = \mathcal{N}(z^*, \sigma_{z^*})$ where $\sigma_{\hat{z}^*}$ is the standard deviation in the 'calculated' upper bound that we can vary. Data points larger than the 'calculated' upper bound would have zero likelihood in case I, but to mitigate this we assume that if there is a higher value in the sampled dataset, $\vec{z}$, we replace the 'calculated' upper bound with the 'empirical' upper bound, $\hat{z}^{*\prime} = \max(\hat{z}^*, \max(\vec{z}))$. We can then use this adjusted upper bound $\hat{z}^{*\prime}$ in case I, sampling a new $\hat{z}^*$ each time we resample the dataset $\vec{z}$, $N_r = 600$ times. These adjustments do not effect case II.

## 3. Results

In this section, we implement our framework to calculate the potential height of a storm surge at a point along the coast of New Orleans. We investigate how the maximum storm surge changes as we vary the parameters of a tropical cyclone. We choose to use the years 2025 to represent the present day climate conditions and 2097, which produces a substantially larger potential intensity and potential size, to represent climate conditions at the end of the century. We use the SSP-585 emissions scenario CESM2-r10i1p1f1 ensemble member August mean to calculate the potential size and potential intensity of the storm. We assume that the tropical cyclone reaches these potentials, i.e maximum intensity and size that are physically plausible, given the monthly average meteorological conditions. We implement the tropical cyclone profile of Chavas *et al.* (2015) to model the respective wind field, and use it to force ADCIRC, varying track displacement, angle, and velocity ($c$, $\chi$, $V_t$).

We choose to use the triangular vertex in the ADCIRC mesh closest to the center of New Orleans for our analysis. We first perform 25 ADCIRC simulations using a Latin hypercube to sample the full range of tropical cyclone parameters. These data points including the simulated surges are used to create an initial fit for the emulator. We then use the MES acquisition function to optimally sample another set of tropical cyclone parameters, and perform another 25 ADCIRC simulations. Figure 7 shows the results of this process, where panels (b-d) describe the tropical cyclone characteristics used to force the $i$-th ADCIRC simulation. Figure 7(a) shows the maximum storm surge height produced from each sample, and the largest heights produced from the full optimization procedure is denoted by the solid line. For this example location, the maximum storm surge attained by this optimization is 0.87 m higher in 2097 than in 2025. There is some uncertainty associated with the optimization procedure, as it is possible that the function has not found the global optima in each case. In Figure 7(b)-(d) we can see that the optimal tropical cyclone track parameters are quite close to one another for the years 2025, as we might expect. Tropical cyclone track displacement is measured east and west of the observation point.

In Figure 9 we show the results from conducting two sets of Bayesian optimisation experiments. We use the CLE15 wind and pressure profiles with the potential size and intensity calculated for August 2025 and August 2097 using the CESM2-r4i1p1f1 SSP585 ensemble member. The locations of the points are shown in Figure 8. The potential height values found for 2097 in (b) are uniformly higher than those for 2025 in (a) as shown in panel (c), as expected because the potential height and potential intensity are higher. In each experiment for all points, the optimal displacement ends up being slightly
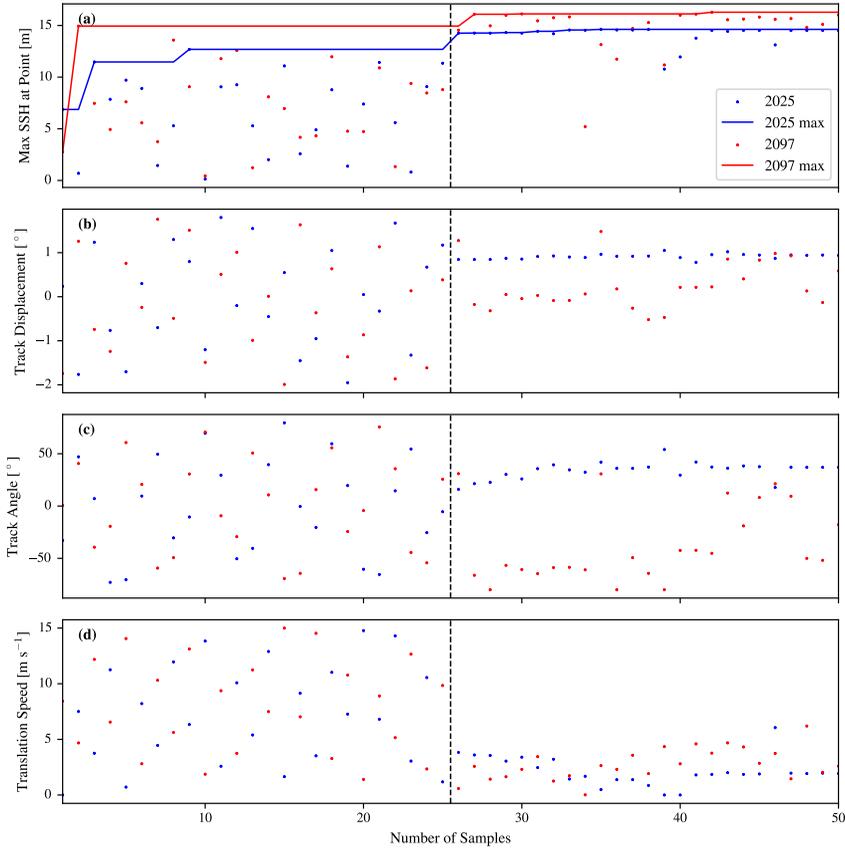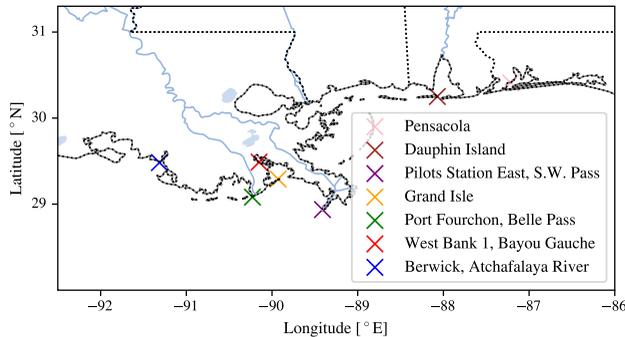
**Figure 7.** *Panel a shows the maximum storm surge height at the point nearest to New Orleans in the mesh. The colors in the figure correspond to potential size and intensity calculated based on 2025 (blue) and 2097 (red) in SSP-585 CESM2-r10i1p1f1 August monthly average. The parameters chosen for each sample are shown in panels b-d. The first 25 points are chosen by Latin hypercube sampling, and this can be seen by the panels b-d filling the whole parameter space, and the second 25 points are the points chosen by the MES acquisition function. The solid lines in panel (a) dynamically show the maximum storm surge height reached in the optimization procedure.*

to the west of the observation point, which would likely correspond to the tropical cyclone's top right corner passing over the points, which agrees with our expectations. The optimal parameters are in general similar between the two years for both locations for every point (panel c), apart from for the point near Dauphin Island, where the track displacement, track angle, and translation speed are all drastically different between the two, despite it following the general trend of an increased Max SSH. This could be caused either by the fact that Bayesian optimization can find local maxima rather than the global maximum, or by the fact that if there are many similar local maxima within the optimization space, a small change in the height of each could lead to a drastic change in the optimal parameters of the global maximum. For the other points, the significant change is the movement of the track displacement slightly westward (negatively in displacement), which could correspond to an increase in the potential size $r_a$ and radius of maximum winds $r_{max}$ between the two years tested. The two headlands of Port Fourchon and Pilots Station East both have optima in both years with the maximum allowed translation speed of 15 m s$^{-1}$ whereas the more enclosed points at Pensacola, West Bank 1, and Berwick all have much slower optimal translation speeds (1 m s$^{-1}$, 7.5 m s$^{-1}$, & 2 m s$^{-1}$ respectively). This is compatable with

Lockwood *et al.* (2022)'s finding that open coastal points have higher storm surge heights at higher TC translation speeds, whereas semi-enclosed coastal points have lower storm surge heights at higher TC translation speeds. The average maximum SSH (potential height) found for the set of points increases by 0.8 m or 11% between 2025 and 2097, equivalent to an average rise in potential height of around 1 cm yr$^{-1}$.



*Figure 8. Reference points chosen to apply the Bayesian optimization procedure to, from the simplified ADCIRC grid used for the New Orleans region. They were chosen to be the closest node in the mesh to the location of the tide gauge station included as the label.*

To highlight the usefulness of these estimates, we run a statistical simulation to investigate the impact of implementing a determined upper bound on the sampling uncertainty given $N_s$ years of block maxima observations at the point of interest. To do this, we fit the generalized extreme value distribution both with and without assuming the upper bound we have computed, i.e. the worst case storm surge height. We assume that the maximum storm surge height observed each year is sampled from the "true" generalized extreme value (GEV) distribution. In Figure 10(a) a sample of 50 data points ($N_s = 50$) from the assumed GEV ($\alpha = 2$ m, $\beta = 1$ m, $\gamma = -0.2$) are denoted by the black dots. Their return period $p$ is estimated as $p = (N_s + 1)/(r)$ where $N_s$ is the number of yearly maxima sampled and $r$ is the rank (order) of the sample, i.e. $\{1, ..., N_s\}$ in descending order. For the GEV distribution itself, the return period of the distributions $p$ for a return value $v$ is calculated as $p = 1/(1 - F(v))$ where $F$ is the cumulative distribution of the fitted distribution. We fit a GEV to this set of sampled data points, assuming the upper bound ($z^* = 7$ m) is known (I, green) or unknown (II, orange). We see that (I) is very close to the true distribution (black), whereas (II) is substantially different. The dashed lines show the upper bound of the fitted distribution for the case where the upper bound was known ahead of time (green) and the case where it was not (orange). For this example the fit for case I (green) where the upper bound was known ahead of time is much closer to the true distribution (black line) than case II (orange), where the upper bound was not known. In particular, case II significantly underpredicts the upper bound compared to the true value, whereas case I is guaranteed to have the same upper bound.

To explore this effect more systemically, we repeat the implementation $N_r = 600$ times for various numbers of samples ($N_s$). The number of samples taken ($N_s$) is varied between $N_{s\ min} = 15$ and $N_{s\ max} = 1000$, and for each number of samples the original GEV is resampled $N_r = 600$ times with different random seeds. The distributions are refitted for each of the $N_r = 600$ resamples, and the 1 in 100 year and 1 in 500 year return values of that fit are calculated. From the $N_r = 600$ resamples for each sample size $N_s$ we calculate the mean fit prediction over the resamples (solid lines) and the 5% to 95% envelope (colored envelopes). This is shown in Figure 10(b)&(c), where we can see that a mean prediction across the resamples in either case is very close to the true value (black line) for both the case where the upper
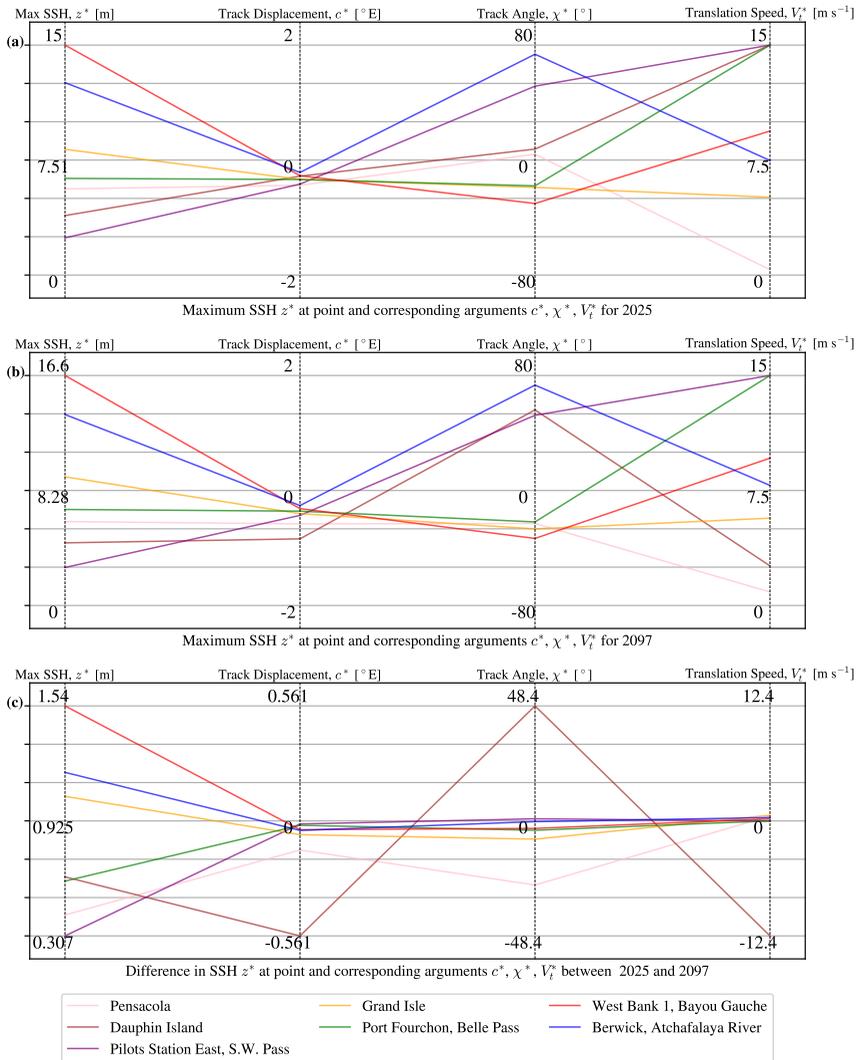
**Figure 9.** *Parallel coordinate plots for the optimal SSH $z^*$ and corresponding optimal TC characteristics ($c^*$, $\chi^*$, $V_t^*$) found for a variety of points near New Orleans (see Figure 8) using potential intensity and size calculated from the CESM2-r4i1p1f1 ensemble member for (a) 2025 (b) 2097 and (c) the difference between the two years. (a) and (b) are plotted between the range imposed as constraints on the optimization.*

bound is known (green line), and unknown (orange line) for the 1 in 100 and 1 in 500 year return values independent of the number of samples taken, $N_s$. However, as expected, the confidence envelope for both substantially declines as the number of samples, $N_s$, is increased. For both return levels, the effect of knowing the upper bound (green envelope) substantially reduces the 5th-95th percentile envelopes when compared to not knowing the upper bound (orange envelope). This effect is more substantial at the 1 in 500 year return level than the 1 in 100 year return level, and for example at $N_s = 51$, the 5%–95% range is reduced 1.98× for the 1 in 100 year return value, and 3.18× for the 1 in 500 year value.

    To investigate whether the results still hold when we vary the uncertainty in the 'calculated' upper bound, we vary the assumed Gaussian noise level while keeping the number of samples at $N_s = 50$ in
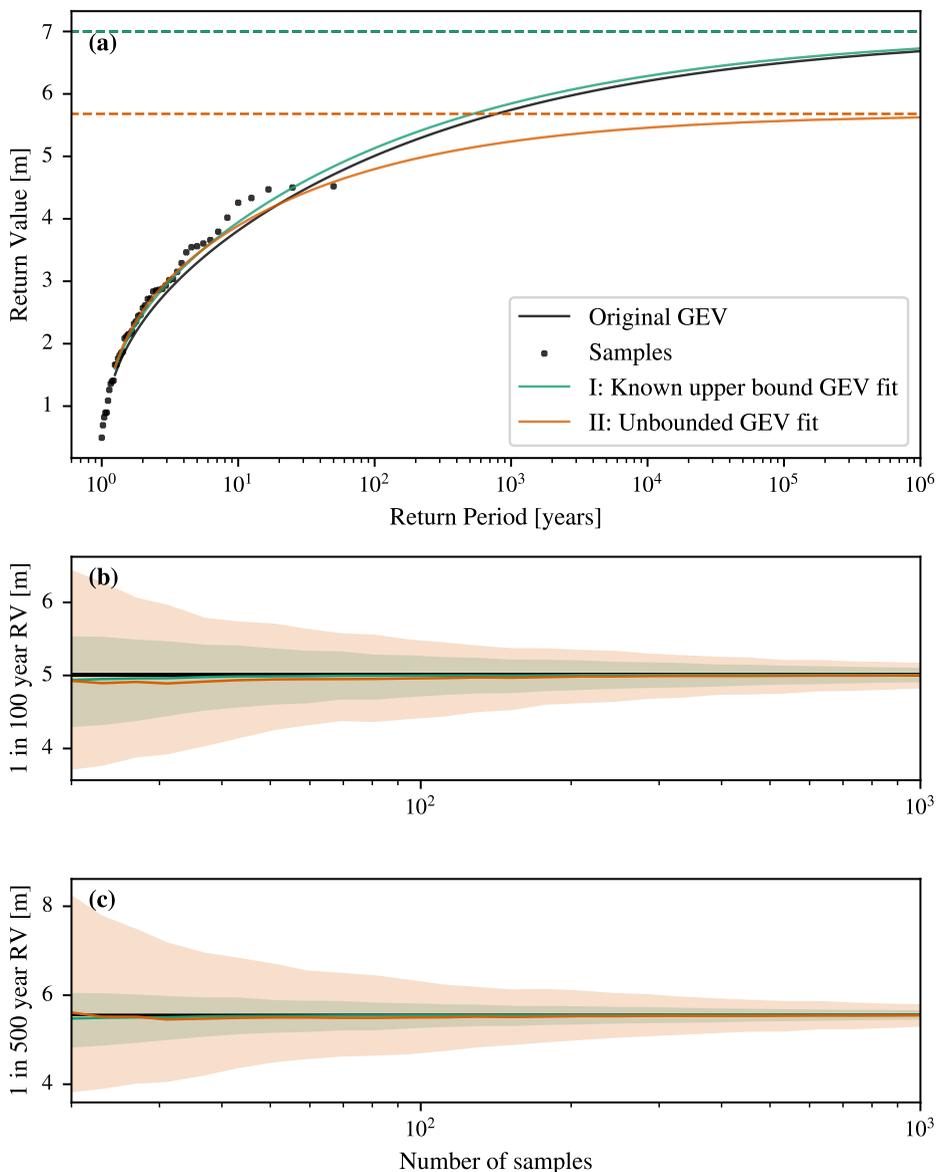
***Figure 10.*** *We investigate how having an upper bound changes the estimated annual exceedance probability (AEP) at selected tide gauge stations. Panel (a) shows a single experiment of sampling 50 years of data (black crosses) from an original GEV (black line) with plausible parameters ($\alpha = 2\,m, \beta = 1\,m, \gamma = -0.2$), and then fitting return periods with the upper bound known (green line) or not known (orange line). Knowing the upper bound ahead of time improves the estimate of the 1 in 100 year and 1 in 500 year event. The solid green and orange lines in each line correspond to the mean prediction from each fitting method over the $N_r = 600$ resamples with different random seed random seeds, and the green and orange envelopes correspond to the area between the 5th percentile and 95th percentile estimate. In panels (b) and (c) the unbounded GEV fit (II) has a much larger 5%–95% envelope (orange) than for the bounded GEV fit (I) (green).*
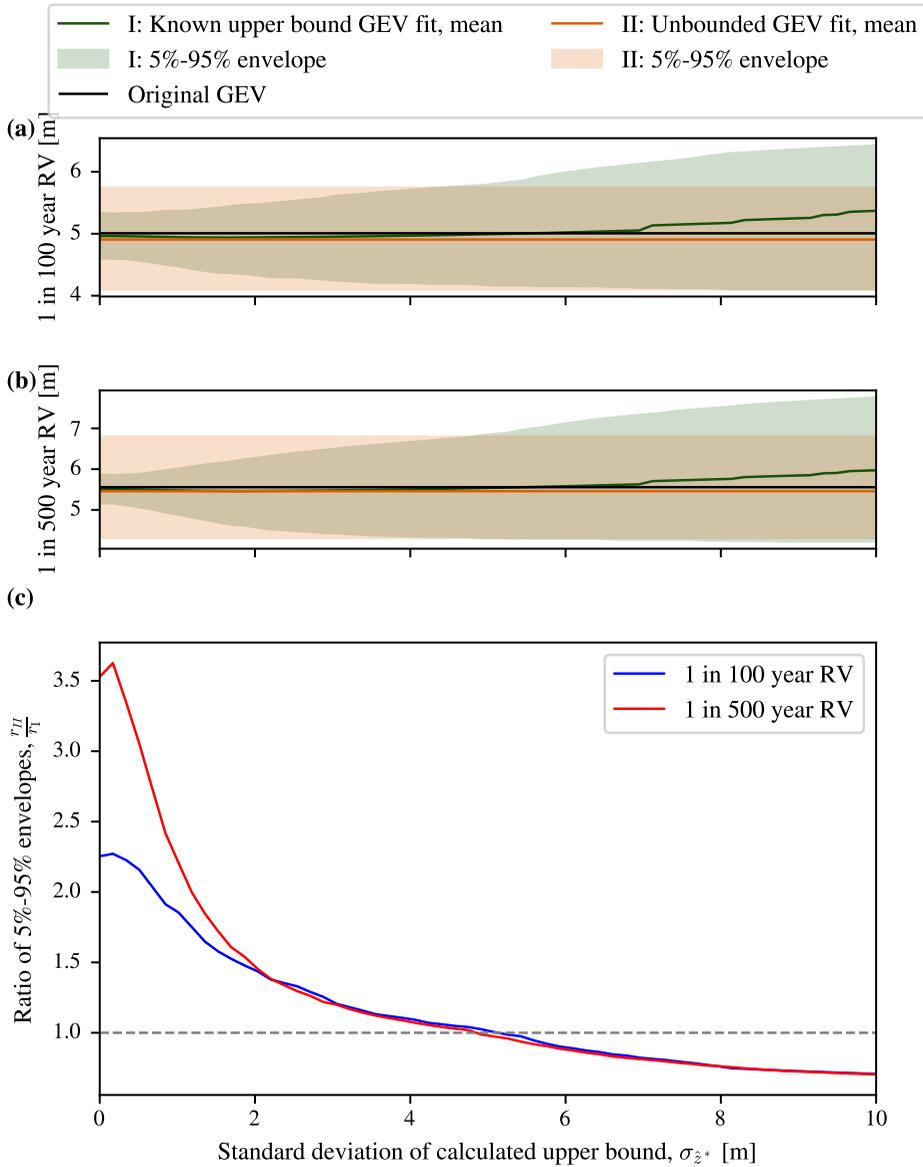
***Figure 11.*** *An experiment into the effect of uncertainty in the upper bound on its usefulness in improving our estimation of high return values. Panels (a) and (b) show the 1 in 100 and 1 in 500 year return values respectively for the original GEV (black line, $z^* = 7\,m, \beta = 1\,m, \gamma = -0.2$). They also show the mean, 5% and 95% intervals calculated by resampling $N_r = 600$ times, with $N_s = 50$ samples each time, and refitting with model fit methods (I) and (II). Fitting method (II) does not depend on $\sigma_{\hat{z}^*}$, and so this is only conducted once. The solid lines in panels are (a) and (b) are the means and the shaded areas are the 5% to 95% model envelopes. Panel (c) summarises the relative evolution of the 5% to 95% range of fitting method (II), $r_{II}$ divided by fitting method (I) $r_{II}$, $\frac{r_{II}}{r_I}$, as $\sigma_{\hat{z}^*}$ increases. This shows that the advantage of method (I) is initially large at small $\sigma_{\hat{z}^*}$, but declines with greater $\sigma_{\hat{z}^*}$, before it becomes disadvantageous at $\sigma_{\hat{z}^*} > 5\,m$ for both return values .*

Figure 11. Figures 11 (a) & (b) show the mean and the 5%-95% envelopes for the 1 in 100 and 1 in 500 year return values (RV) for fitting methods (I) and (II), where we change the standard deviation in the calculated upper bound $\sigma_{\hat{z}^*}$ for method (I). As shown in the two panels, an estimate of the upper bound (I) reduces the bias in the mean estimate, and leads to a smaller 5%–95% envelope, until is very large $\sigma_{\hat{z}^*} > 5$ m. As $\sigma_{\hat{z}^*}$ progresses from 0 m to 5 m (where the scale is 1 m), the range of (I) approaches (II). To show this quantitatively, in Figure 11(c) we can see that the ratio between the 5%-95% confidence envelopes begins larger for the 1 in 500 year events than the 1 in 100 year events, as knowing the upper bound has a larger effect on our estimation of longer return period return values as before. Both reach 1.0 at $\sigma_{\hat{z}^*} \approx 5$ m as the standard deviation in the calculated upper bound increases and the information becomes less informative. This might be surprising because it suggests that knowing the upper bound in this set up, even with this large standard deviation, improves the return value estimate. This could be because if the calculated upper bound is too low, then it is likely to be replaced by the maximum of the samples, and if it is too high, then the shape and scale parameters can change to compensate and achieve good long period RV, but this should be studied further. When we go beyond this, $\sigma_{\hat{z}^*} > 5$ m, then the range of method (I) is greater than method (II) as shown in Figure 11(c). We can also see that the mean prediction in Figure 11 (a) & (b) begins to increase sharply at $\sigma_{\hat{z}^*} > 7$ m, suggesting that at this point the unrealistically high upper bounds assumed do have a negative impact. The bias is positive because if an unrealistically low upper bound is sampled for $\hat{z}^*$ then it is likely to be replaced with the maximum of the observations max $(\vec{z})$.

## 4. Discussion

### 4.1. Summary of Findings

Through our initial results we have shown the implications of finding the maximum potential height of a tropical cyclone storm surge and how it may increase due to climate change. Key findings of our work include: (i) that Bayesian optimization can be used to more efficiently find the maximum storm surge height produced by a tropical cyclone with a variable trajectory, (ii) that the combined increase in potential size and potential intensity leads to a mean increase of 8% in the maximum storm surge height at seven points along the coast of New Orleans over the 21st century, and (iii) that in idealized simulations, knowledge of the upper bound of the potential storm surge height can substantially reduce the uncertainty in the 1 in 100 and 1 in 500 year return values, even if there is uncertainty in the upper bound. It is our hope that this work will provide foundational bases for further work on the topic.

### 4.2. Potential Size

sec:potential$_size$

To our knowledge, this is the first time that potential size has been calculated for a specific geographic region or as a time-varying characteristic. It is also the first time we have seen it utilized in constraining estimates of storm surge risk. By driving the ADCIRC model with the CLE15 profile, and assuming an isothermal inflow when integrating to find the pressure profile, we made our forcing method as consistent with the potential size calculation as possible. Our implementation allows us to show that the potential size is expected to increase in the future and is highly correlated to the sea surface temperature. We found that the potential size increases by around 2.3 km per year in the SSP585 ensemble members that we investigated.

Here we have used just a single CMIP6 model member and a single high-emission scenario to exemplify the use of our framework for investigating possible effects of climate change on the storm surge hazard. SSP-585 is now considered to be a less likely emissions scenario (e.g. Huard *et al.* (2022)), and so more likely results for risk quantification can be explored by using intermediate scenarios. It would be worthwhile to compare the results produced by using different CMIP6 models and scenarios

to explore more systematically whether different ensemble members exhibit different relationships to potential intensity and size.

We also use monthly averages to calculate potential intensity and size, but potential intensity and size could vary significantly within months. For example, marine heat waves could increase the potential intensity of tropical cyclones. An alternative approach is to consider a particular percentile of potential intensity, as was done in Mori *et al.* (2022), to consider, for example, the 1 in 10 year potential intensity, and work out the corresponding storm surge that this would produce.

There are some key issues that might affect our estimates of potential size and intensity. First, the potential size measure does not include a number of important processes, such as the dissipation of energy through eddy shedding (Wang *et al.* 2022a). As noted by Wang *et al.* (2022a), their derivation also excludes some time-varying effects, such as the ventilation of colder water to the surface, which is also true for potential intensity as calculated (Bister & Emanuel 2002). Additionally, it relies on a number of tuneable parameters, such as the dissipation rate, $w_{\mathrm{cool}}$, the efficiency relative to the Carnot engine, $\eta_{\mathrm{rel.carnot}}$, and the supergradient factor, $\gamma_{\mathrm{sg}}$. As these values may have been tuned to match the numerical simulations of azimuthally symmetric TCs in Wang *et al.* (2022a), the potential size has some degree of subjectivity. Wang *et al.* (2022a) compares the calculated potential size against numerical simulations of tropical cyclones, but they do not compare the measure against observations. Comparing tropical cyclone sizes recorded in terms of $r_{\mathrm{max}}$ from the International Best Track Archive for Climate Stewardship (IBTrACS, Knapp *et al.* (2018)) to the potential sizes calculated using ERA5 data (Hersbach *et al.* 2020) would support verification of the potential size model, and allow us to assess the distribution of sizes as a fraction of potential size that has been observed (similar to Emanuel (2000)).

Finally, there may be broader theoretical inconsistencies made in the derivations of potential intensity and size that we have not explored here. First, though Emanuel (1986)'s potential intensity is well-established and relatively well accepted (Rousseau-Rizzi *et al.* 2021), Makarieva & Nefiodov (2023) suggests that the model has an internal thermodynamic inconsistency. There are also a number of assumptions made in the Chavas *et al.* (2015) profile that join together the analytic solutions from Emanuel & Rotunno (2011) and Emanuel (2004), both of which were analytically derived. Wang *et al.* (2022a) and Wang *et al.* (2023) use a more detailed Carnot engine to calculate the pressure drop to the radius of maximum wind than Emanuel (1986), and so recalculating the potential intensity $V_p$ in a way that is congruent with this calculation is worth further exploration.

### 4.3. Bayesian optimization

We have shown that Bayesian optimization can be used to find the potential height at different points along the coast, for two different years, showing the increase in potential height between them. For the set of seven points near New Orleans that were chosen, the average increase in calculated potential height was 0.8 m (11%) between August 2025 and August 2097.

In Figure 5 we show that our framework improves upon randomly selecting points in terms of minimizing simple regret. We used the MES data acquisition function of Wang & Jegelka (2017) to optimize the trajectory of a tropical cyclone, which we expected to outperform the expected improvement acquisition function used in Ide *et al.* (2024). However have not compared our strategy against any other data acquisition functions, or verified that the strategy is superior for all the points along the coast. Bayesian optimization is also dependent on the performance of the emulator in capturing the relationship between the variables in the modelled space, and so adapting the kernel to be more appropriate to the problem (as in Tazi *et al.* (2023)), could lead to the strategy minimizing regret more effectively.

### 4.4. Statistical simulation

We use statistical simulation to investigate the utility of knowing the upper bound in improving the estimation of return values, based on samples taken from a GEV distribution. First, we investigate the

effect of varying the sample size $N_s$ on our estimates of the return values (in Figure 10), and find that knowing the true upper bound $z^*$ improves the estimate of the return values, especially for smaller sample sizes and for the longer return period. Secondly, we assume some Gaussian error in the upper bound to investigate what would happen if we did not know this upper bound exactly, keeping $N_s = 50$ (Figure 11). We find that the uncertainty is improved with a moderate error ($\sigma_{\hat{z}^*} \approx 1$ m) in the potential height of the storm surge, so we could expect it to improve relevant risk estimates, and only if it is extremely imprecise (e.g. $\sigma_{\hat{z}^*} > 5$ m) would it reduce the quality of risk estimates. For both of these statistical simulations, we do not vary the values of $\beta$ and $\gamma$, but we expect (based on Figure 17) in Appendix 8, that our findings will also generalize to these parameters if they were varied for both experiments. Through these results, we have shown that calculating the potential height could usefully augment observations even with realistic calculation uncertainty.

Nonetheless, a number of caveats and unexplored questions remain. None of these simulations make use of real observations. As such, we have not fully explored the practical implications of applying this technique for different points on the coast or evaluating the change in estimated risk. Additionally, in our work we maximized the log-likelihood of the GEV parameter values given the data samples for simplicity. However, Jewson *et al.* (2025) show that doing so would lead to a negative bias where return periods are exceeded more often than would be expected, and that it would be better to use calibrating priors. Therefore, it is worth exploring whether the effects we see in these statistical experiments would still hold if we use a more advanced set of statistical assumptions. We have still assumed that the observations are i.i.d. and therefore ignore non-stationarity in the observations created by climate change.

### 4.5. Storm surge model improvements

To calculate the potential height of the storm surge we used an implementation of the ADCIRC model. Our implementation does not include all of the processes involved in creating the highest potential extreme water height along the coast. Specifically, we have excluded the effects of tides (including tide-surge interaction), wave run up and set-up, and sea level rise. To estimate the potential total water height that could be generated at a point along the coast, these processes should be included. Future studies can include tides in the ADCIRC model simulations, which can be optimized over by changing the time of impact in the simulation of the tropical cyclone around the spring tide. The effects of waves can be included by using the ADCIRC's popular coupling with the SWAN wave model (Booij *et al.* 1999), though this will add significant computational cost. Sea level rise con be included by adding a constant offset of additional water height to the existing ADCIRC simulation, taken from the same CMIP6 scenarios for self-consistency. Finally, Chaigneau *et al.* (2024) showed that using the baroclinic NEMO model improves the accuracy of simulated storm surge heights inareas such as on the Southeastern Florida Peninsula. Future research could use this model in place of the (barotropic) ADCIRC model, however this also comes with increased computational cost (more than 100 times).

In additional to climate change, the shape and position of the coastline will also change over the next 100 years, likely affecting the vulnerability of different communities to storm surges. For example, many coastal areas are undergoing subsidence (e.g. Nicholls *et al.* (2021)). Representing this in the ADCIRC model would necessitate changing the computational mesh. Different projections of the morphology of the coastline as well as the resulting dynamic sea level rise could be incorporated in the model in future studies.

### 4.6. Generalizations of this framework

The framework we have developed can be extended to any coastline and potentially to other environmental hazards such as extreme rainfall flooding, e.g. if comparable limits for rainfall were suggested. Here, we have focused on the US Gulf of Mexico and the New Orleans area in particular, but we could extend this work to more areas, such as East Asia. Indeed, the areas of the largest growth in economic

exposure are expected in East, South East, and South Asia (e.g. for the Pearl River Delta, Deng *et al.* (2022)), due to population growth and economic growth in the coastal regions of these areas.

Both the total height of the water and the duration of its elevation above a certain level can have an impact on the damage done to a coastal community, as these provide more time for flood defences to fail, for flood water to propagate inland, and for water from pluvial and river sources to compound the flooding. While it may be much more difficult to construct physical bounds for tropical cyclone rainfall, and for the pre-existing river levels that respond to that rainfall, it is worth considering optimization of a more combined measure that represents the damage done to the coastline such as that described in e.g. Zhang *et al.* (2000).

## 5. Conclusion

We have seen that Bayesian optimisation can be used to efficiently relate the potential intensity and size of a tropical cyclone to the corresponding maximum potential height of a storm surge at a point along the coast. Both potential intensity and size generally increase assuming projected climate change scenarios, and we are able to use the developed framework to estimate the corresponding increase in the potential height of associated storm surges. We have shown that knowledge of the maximum potential storm surge height can effectively reduce the uncertainty in return periods important for practical applications. However, we do not make use of data produced by nearby coastal points or different forcing scenarios, and this framework could be usefully extended by doing so. Further, we do not fully explore bias correction of climate data, which is important to understanding how tropical cyclone potential size, intensity, and storm surge height change over time. We have developed our framework and made it available as a set of open-source Python packages, hopefully enabling the wider community to easily build upon this work. In the future, we hope this framework can be used to provide more robust and accurate risk estimation for different areas of the world.

## Conflict of interest

The authors declare no conflict(s) of interest.

## Data availability statement

The GEV simulation data plotted was lightweight so has been added to the code repository (Thomas 2025). The processed CMIP6 data used was too bulky to easily share, but the raw data is easily available from the Pangeo catalogue (Pangeo 2022).

## Open source software

The main code to produce the figures for this paper is available at https://github.com/sdat2/worstsurge (Thomas 2025), with readthdocs documentation at https://worstsurge.readthedocs.io/en/latest/MAIN_README.html, and builds on some utility scripts in the `sithom` package (Thomas 2024). Bayesian optimization is provided by the `trieste` package (Picheny *et al.* 2023) which uses `gpflow` (Matthews *et al.* 2017) and `tensorflow` (Abadi *et al.* 2015). ADCIRC is publicly available at https://github.com/adcirc/adcirc, and our slightly modified copy for the Archer2 HPC is available at https://github.com/sdat2/adcirc-v55.02 with our compilation settings. The CLE15 tropical cyclone profile of (Chavas *et al.* 2015) was calculated using their original `matlab` implementation (Chavas 2022), but using `octave` for accessibility.

# References

1. Danabasoglu, G. *et al.* The community earth system model version 2 (CESM2). *Journal of Advances in Modeling Earth Systems* **12,** e2019MS001916 (2020).

2. Camargo, S. J. & Wing, A. A. Tropical cyclones in climate models. *Wiley Interdisciplinary Reviews: Climate Change* **7,** 211–237 (2016).

3. Sainsbury, E. M. *et al.* Can low-resolution CMIP6 ScenarioMIP models provide insight into future European post-tropical-cyclone risk? *Weather and Climate Dynamics* **3,** 1359–1379 (2022).

4. Peng, Y. & Guo, Y.-P. Seasonal cycle delay of the western North Pacific tropical cyclone genesis frequency in CMIP6 simulations. *Geophysical Research Letters* **51,** e2024GL109150 (2024).

5. Shan, K., Lin, Y., Chu, P.-S., Yu, X. & Song, F. Seasonal advance of intense tropical cyclones in a warming climate. *Nature* **623,** 83–89 (2023).

6. Roberts, M. J. *et al.* Projected future changes in tropical cyclones using the CMIP6 HighResMIP multimodel ensemble. *Geophysical research letters* **47,** e2020GL088662 (2020).

7. Sobel, A. H. *et al.* Near-term tropical cyclone risk and coupled Earth system model biases. *Proceedings of the National Academy of Sciences* **120,** e2209631120 (2023).

8. Emanuel, K. A. The Theory of Hurricanes. *Annual Review of Fluid Mechanics* **23,** 179–196 (1991).

9. Emanuel, K. A. An air-sea interaction theory for tropical cyclones. Part I: Steady-state maintenance. *Journal of the Atmospheric Sciences* **43,** 585–605 (1986).

10. Emanuel, K. Tropical cyclones. *Annual review of earth and planetary sciences* **31,** 75–104 (2003).

11. Emanuel, K. A. Hurricanes: Tempests in a greenhouse. *Physics Today* **59,** 74–75 (2006).

12. Rousseau-Rizzi, R., Rotunno, R. & Bryan, G. A thermodynamic perspective on steady-state tropical cyclones. *Journal of the Atmospheric Sciences* **78,** 583–593 (2021).

13. Wehner, M. F. & Kossin, J. P. The growing inadequacy of an open-ended Saffir–Simpson hurricane wind scale in a warming world. *Proceedings of the National Academy of Sciences* **121,** e2308901121 (2024).

14. Mori, S., Shimura, T., Miyashita, T., Webb, A. & Mori, N. Future changes in extreme storm surge based on a maximum potential storm surge model for East Asia. *Coastal Engineering Journal* **64,** 630–647 (2022).

15. Wang, D., Lin, Y. & Chavas, D. R. Tropical cyclone potential size. *Journal of the Atmospheric Sciences* **79,** 3001–3025 (2022).

16. Luettich Jr, R. A. & Westerink, J. J. A solution for the vertical variation of stress, rather than velocity, in a three-dimensional circulation model. *International Journal for Numerical Methods in Fluids* **12,** 911–928 (1991).

17. Jia, G. & Taflanidis, A. A. Kriging metamodeling for approximation of high-dimensional wave and surge responses in real-time storm/hurricane risk assessment. *Computer Methods in Applied Mechanics and Engineering* **261,** 24–38 (2013).

18. Ide, Y., Ozaki, S., Yamashiro, M. & Kodama, M. Development and improvement of a method for determining the worst-case typhoon path for storm surge deviation through Bayesian optimization. *Engineering Applications of Artificial Intelligence* **132,** 107950 (2024).

19. Irish, J. L., Resio, D. T. & Cialone, M. A. A surge response function approach to coastal hazard assessment. Part 2: Quantification of spatial attributes of response functions. *Natural hazards* **51,** 183–205 (2009).

20. Lin, N. & Emanuel, K. A. Grey swan tropical cyclones. *Nature Climate Change 2015 6:1* **6,** 106–111. ISSN: 1758-6798 (Aug. 2015).

21. Schwerdt, R. W., Ho, F. P. & Watkins, R. R. Meteorological criteria for standard project hurricane and probable maximum hurricane windfields, Gulf and East Coasts of the United States. *NOAA technical Report NWS 23, National Oceanic and Atmospheric Administration, US Department of Commerce, Washington, D.C.* (1979).

22. ONR, U. Safety Assessment Principles for Nuclear Facilities. *Bootle: ONR* (2014).

23. Wang, D., Lin, Y. & Chavas, D. R. Corrigendum to "Tropical cyclone potential size". *Journal of the Atmospheric Sciences* **80,** 1657–1660 (2023).

24. Chavas, D. R., Lin, N. & Emanuel, K. A. A Model for the Complete Radial Structure of the Tropical Cyclone Wind Field. Part I: Comparison with Observed Structure. EN. *Journal of the Atmospheric Sciences* **72.** Publisher: American Meteorological Society Section: Journal of the Atmospheric Sciences, 3647–3662. ISSN: 0022-4928, 1520-0469 (Sept. 2015).

25. Wang, Z. & Jegelka, S. *Max-value entropy search for efficient Bayesian optimization* in *International Conference on Machine Learning* (2017), 3627–3635.

26. Huard, D., Fyke, J., Capellán-Pérez, I., Matthews, H. D. & Partanen, A.-I. Estimating the likelihood of GHG concentration scenarios from probabilistic Integrated Assessment Model simulations. *Earth's Future* **10,** e2022EF002715 (2022).

27. Pangeo. *Pangeo Data Catalog* 2022. https://pangeo.io/catalog.html (June 20, 2022).

28. Busecke, J., Ritschel, M., Maroon, E. & Nicholas, T. *jbusecke/xMIP: v0.7.1* version v0.7.1. Jan. 2023. doi:10.5281/zenodo.7519179. https://doi.org/10.5281/zenodo.7519179.

29. Schulzweida, U. *CDO User Guide* version 2.3.0. Oct. 2023. doi:10.5281/zenodo.10020800. https://doi.org/10.5281/zenodo.10020800.

30. Bister, M. & Emanuel, K. A. Low frequency variability of tropical cyclone potential intensity 1. Interannual to interdecadal variability. *Journal of Geophysical Research: Atmospheres* **107,** ACL 26–1–ACL 26–15 (2002).

31. Gilford, D. M. pyPI (v1. 3): Tropical cyclone potential intensity calculations in Python. *Geoscientific Model Development* **14,** 2351–2369 (2021).

32. Chavas, D. R., Reed, K. A. & Knaff, J. A. Physical understanding of the tropical cyclone wind-pressure relationship. *Nature communications* **8,** 1360 (2017).

33. Luettich, R. A., Westerink, J. J., Scheffner, N. W., *et al.* ADCIRC: an advanced three-dimensional circulation model for shelves, coasts, and estuaries. Report 1, Theory and methodology of ADCIRC-2DD1 and ADCIRC-3DL (1992).

34. Westerink, J. J., Luettich Jr, R., Blain, C. & Scheffner, N. W. *Adcirc: an advanced three-dimensional circulation model for shelves, coasts, and estuaries. report 2. user's manual for adcirc-2ddi* tech. rep. (Army Engineer Waterways Experiment Station Vicksburg Ms, 1994).

35. Dietrich, J. *et al. Real-time forecasting and visualization of hurricane waves and storm surge using SWAN+ ADCIRC and FigureGen* in *Computational challenges in the geosciences* (2013), 49–70.

36. Garnett, R. *Bayesian optimization* (Cambridge University Press, 2023).

37. Khatamsaz, D. *et al.* A physics informed bayesian optimization approach for material design: application to NiTi shape memory alloys. *npj Computational Materials* **9,** 221 (2023).

38. Hashemi, M. R., Spaulding, M. L., Shaw, A., Farhadi, H. & Lewis, M. An efficient artificial intelligence model for prediction of tropical storm surge. *Natural Hazards* **82,** 471–491 (2016).

39. Williams, C. K. & Rasmussen, C. E. *Gaussian processes for machine learning* **3** (MIT press Cambridge, MA, 2006).

40. Gopinathan, D., Heidarzadeh, M. & Guillas, S. Probabilistic quantification of tsunami current hazard using statistical emulation. *Proceedings of the Royal Society A* **477,** 20210180 (2021).

41. McKay, M. D., Beckman, R. J. & Conover, W. J. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21,** 239–245 (1979).

42. Tazi, K. *et al.* Beyond intuition, a framework for applying GPs to real-world data. *arXiv preprint arXiv:2307.03093* (2023).

43. Coles, S., Bawa, J., Trenner, L. & Dorazio, P. *An introduction to statistical modeling of extreme values* (Springer, 2001).

44. Kingma, D. P. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

45. Lockwood, J. W., Lin, N., Oppenheimer, M. & Lai, C.-Y. Using Neural Networks to Predict Hurricane Storm Surge and to Assess the Sensitivity of Surge to Storm Characteristics. *Journal of Geophysical Research: Atmospheres,* e2022JD037617 (2022).

46. Knapp, K. R., Diamond, H. J., Kossin, J. P., Kruk, M. C., Schreck, C., *et al.* International best track archive for climate stewardship (IBTrACS) project, version 4. *NOAA National Centers for Environmental Information* (2018).

47. Hersbach, H. *et al.* The ERA5 global reanalysis. en. *Quarterly Journal of the Royal Meteorological Society* **146,** 1999–2049. ISSN: 1477-870X (2020).

48. Emanuel, K. A statistical analysis of tropical cyclone intensity. *Monthly weather review* **128,** 1139–1152 (2000).

49. Makarieva, A. M. & Nefiodov, A. V. A critical analysis of the assumptions underlying the formulation of maximum potential intensity for tropical cyclones. *Journal of the Atmospheric Sciences* **80,** 1201–1209 (2023).

50. Emanuel, K. & Rotunno, R. Self-stratification of tropical cyclone outflow. Part I: Implications for storm structure. *Journal of the Atmospheric Sciences* **68,** 2236–2249 (2011).

51. Emanuel, K. A. Tropical cyclone energetics and structure. *Atmospheric turbulence and mesoscale meteorology* **165,** 192 (2004).

52. Jewson, S., Sweeting, T. & Jewson, L. Reducing reliability bias in assessments of extreme weather risk using calibrating priors. *Advances in Statistical Climatology, Meteorology and Oceanography* **11,** 1–22 (2025).

53. Booij, N., Ris, R. C. & Holthuijsen, L. H. A third-generation wave model for coastal regions: 1. Model description and validation. *Journal of geophysical research: Oceans* **104,** 7649–7666 (1999).

54. Chaigneau, A. A., Menéndez, M., Ramírez-Pérez, M. & Toimil, A. Regional modelling of extreme sea levels induced by hurricanes. *Natural Hazards and Earth System Sciences Discussions* **2024,** 1–29 (2024).

55. Nicholls, R. J. *et al.* A global analysis of subsidence, relative sea-level change and coastal flood exposure. *Nature Climate Change* **11,** 338–342 (2021).

56. Deng, Z., Wang, Z., Wu, X., Lai, C. & Zeng, Z. Strengthened tropical cyclones and higher flood risk under compound effect of climate change and urbanization across China's Greater Bay Area. *Urban Climate* **44,** 101224 (2022).

57. Zhang, K., Douglas, B. C. & Leatherman, S. P. Twentieth-century storm activity along the US east coast. *Journal of Climate* **13,** 1748–1761 (2000).

58. Thomas, S. D. A. *WorstSurge - Finding the potential height of a tropical cyclone in a changing climate using Bayesian Optimization* version v0.0.1. Mar. 2025. doi:10.5281/zenodo.15096627. https://github.com/sdat2/worstsurge.

59. Thomas, S. D. A. *Sithom's Scientific Python Utilities Package* version v0.1.1. Oct. 2024. doi:10.5281/zenodo.7020109. https://github.com/sdat2/sithom.

60. Picheny, V. *et al.* Trieste: Efficiently exploring the depths of black-box functions with tensorflow. *arXiv preprint arXiv:2302.08436* (2023).

61. Matthews, A. G.d. G. *et al.* GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research* **18,** 1–6 (2017).

62. Abadi, M. *et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems* Software available from tensorflow.org. 2015. https://www.tensorflow.org/.

63. Chavas, D. R. *Code for tropical cyclone wind profile model of Chavas et al (2015, JAS)* 2022. doi:doi://10.4231/CZ4P-D448. https://purr.purdue.edu/publications/4066/1.

64. Wang, S., Lin, N. & Gori, A. Investigation of tropical cyclone wind models with application to storm tide simulations. *Journal of Geophysical Research: Atmospheres* **127,** e2021JD036359 (2022).

65. Lin, N. & Chavas, D. On hurricane parametric wind and applications in storm surge modeling. *Journal of Geophysical Research: Atmospheres* **117** (2012).

66. Holland, G. J., Belanger, J. I. & Fritz, A. A Revised Model for Radial Profiles of Hurricane Winds. *Monthly Weather Review* **138,** 4393–4401 (2010).

## 6. Tropical cyclone potential intensity and potential size details

Figure 12 shows how the potential size is calculated from the environmental variables taken from the CMIP6 model member. The key inputs to first calculate the potential intensity using tcPyPI (Gilford 2021) are the sea surface temperature of the water (from the ocean monthly average table, Omon, TS, $T_s$), the mean sea level pressure (MSL, $p_s$), air temperature column (TA, $T_a(p)$) and specific humidity (Q, $q(p)$) column (all from the atmospheric monthly average table, Amon). The outflow temperature $T_0$ and level $OTL$ is calculated by following a moist adiabat from the sea surface to the point of intersection with the observed atmospheric profile. tcPyPI then calculates the potential intensity, $V_p$, following Bister & Emanuel (2002) as,

$$(V_p)^2 = \frac{T_s}{T_0} \frac{C_k}{C_D} (\text{CAPE}^* - \text{CAPE}) \mid_{\text{RMW}}, \tag{4}$$

where $V_p$ is the potential intensity at the gradient wind level, $\text{CAPE}^*$ is the convective available potential energy of saturated air lifted from the sea surface to the outflow level, and $\text{CAPE}_{\text{env}}$ is the convective potential energy of the environment. The convective available potential energy is conventionally defined as the work done per unit mass by the buoyancy force acting on an air parcel from its level of free convection $h_{\text{LFC}}$ to its level of neutral buoyancy $h_{\text{LNB}}$,

$$\text{CAPE} = \int_{h_{\text{LFC}}}^{h_{\text{LNB}}} B \, dh = g \int_{h_{\text{LFC}}}^{h_{\text{LNB}}} \frac{\rho_e - \rho_p}{\rho_e} \, dh, \tag{5}$$

where $B$ is the Buoyancy force per unit mass, $\rho_p$ is the density of the parcel lifted to $h$ and $\rho_e$ is the density of the environment at $h$, We assume that the air is lifted on a moist adiabat, exchanging no entropy with the surrounding air during ascent. When calculating the potential intensity, the ratio of the surface enthalpy exchange to momentum exchange coefficient $\frac{C_k}{C_D}$ is assumed to be 0.9. Therefore, caling the tcPYPI package ends up looking like,

$$\text{tcPYPI}\left(T_s, p_{\text{msl}}, T_a(p), q(p), \frac{C_k}{C_D}\right) = (V_p, P_{\text{min}}, T_0, \text{OTL}), \tag{6}$$

where $T_s$ is the sea surface temperature, $p_{\text{msl}}$ is the sea level pressure, $T_a(p)$ is the air temperature profile, and $q(p)$ is the specific humidity profile. The minimum central pressure $P_{\text{min}}$ is also calculated, but not used to calculate the potential size $r_a$ or determine the final tropical cyclone to force the ADCIRC model. When calling tcPYPI we do not reduce the potential intensity (feeding $V_{\text{reduc}} = 1$) so that $V_p$ corresponds to the windspeed above the boundary layer rather than the 10m wind to be consistent with the CLE15 model. We later reduce the CLE15 wind profile by $V_{\text{reduc}} = 0.8$ as detailed in Appendix 7 before it is applied to ADCIRC which is the standard value in tcPYPI (Gilford 2021).

To then go from potential intensity to potential size, we use the procedure described in Wang *et al.* (2022a) with some additional assumptions, where we find the outer radius $r_a$ that produces the same

01   pressure at the point of maximum winds $p_m$ with Wang *et al.* (2022a)'s Carnot cycle (W22) and the
02   Chavas *et al.* (2015) tropical cyclone radial atmospheric profile (CLE15).

03   We first assume that the maximum windspeed used in the W22 Carnot engine $V_{\text{max W22}}$ is a supergra-
04   dient constant $\gamma_{\text{sg}} = 1.2$ above the maximum gradient wind assumed for the CLE15 model $V_{\text{gm}}$, which
05   we further assume is the potential intensity calculated by the tcPYPI package, $V_p$ so that

$$V_{\text{max W22}} = \gamma_{\text{sg}} V_{\text{gm}} = \gamma_{\text{sg}} V_p, \tag{7}$$

09   as in Wang *et al.* (2022a) Equation 23. This is allowed in the calculation of potential size (see Wang *et*
10   *al.* (2022a) sections 2b and 2c) but it is not the only possibility. This also introduces some inconsistency
11   between assuming one Carnot cycle from Emanuel (1986) to calculate $V_p$ and a more complex Carnot
12   cycle to calculate the thermodynamic constraint on the central pressure $p_{m2}$, and thereby the potential
13   size $r_a$. Further $\gamma_{\text{sg}}$ could be reasonably set at 1.05 or 1.1, so this introduces additional subjectivity to
14   potential size (Wang *et al.* 2022a).

15   To calculate the potential size, we drive two models; the CLE15 dynamic constraint model and the
16   W22 thermodynamic constraint Carnot engine, both of which provide an estimate of the pressure at
17   the maximum winds, given a number of inputs, and we vary the outer radius of the TC $\tilde{r}_a$ until the
18   two estimates of pressure are within some tolerance, $t$, of one another. The CLE15 model depends on
19   the surface $T_s$ and outflow temperature $T_0$ (calculated from the atmospheric profile using tcPYPI), the
20   background sea level pressure $p_a$, the lower-troposphere subsidence velocity in the subsidence region
21   $w_{\text{cool}} = 0.002 \text{ m s}^{-1} = 2 \times 10^{-3} \text{ m s}^{-1}$, the surface drag coefficient $C_D = 0.0015 = 1.5 \times 10^{-3}$ and the
22   surface enthalpy exchange coefficient $C_k = 0.9 \times C_D = 1.35 \times 10^{-3}$ (to be consistent with potential
23   intensity calculation), the Coriolis parameter $f$ at that latitude, the potential intensity $V_p$ and the outer
24   radius $\tilde{r}_a$ which leads to a prediction of the pressure at the radius of maximum winds $p_{m1}$ and the radius
25   of maximum winds $r_{\text{max}}$,

$$\text{CLE15}\left(V_p, \rho_a, p_a, w_{\text{cool}}, f, C_D, C_k; \tilde{r}_a\right) = \left(p_{m1}, \tilde{r}_{\text{max}}\right). \tag{8}$$

29   That prediction of the pressure $p_{m1}$ is made assuming that the gradient wind of the CLE15 profile,
30   $V(r)$, is in cyclogeostrophic balance, and that the air density is calculated that it is an isothermal ideal
31   gas so that the pressure profile is,

$$p(r) = p_a \exp\left(-\frac{\rho_a}{p_a} \int_r^{\tilde{r}_a} \left(f V(\tilde{r}) + \frac{V^2(\tilde{r})}{\tilde{r}}\right) d\tilde{r}\right), \tag{9}$$

36   and $p_{m1} = p(r_{\text{max}})$. The W22 model again takes the surface and output temperatures $T_s$ and $T_0$, the
37   background sea level pressure $p_a$, the environmental relative humidity $\mathcal{H}_e$, the efficiency relative to the
38   Carnot cycle $\eta_{\text{rel.carnot}} = 0.5$, the lift parametrisation $\beta_l = 1.25$, the Coriolis parameter $f$, the assumed
39   maximum velocity $V_{\text{max W22}} = \gamma_{\text{sg}} V_p$, the radius of maximum winds $\tilde{r}_{\text{max}}$ from CLE15, and the radius
40   of outer winds $\tilde{r}_a$

$$\text{W22}\left(T_s, T_0, p_a, \eta_{\text{rel.carnot}}, f, \mathcal{H}_e, V_{\text{max W22}}, \beta_l; \tilde{r}_{\text{max}}, \tilde{r}_a\right) = p_{m2}. \tag{10}$$

45   To converge on a final value of the outer radius $\tilde{r}_a$ where $|p_{m1} - p_{m2}| < t$ then we can change $\tilde{r}_a$
46   using the bisection algorithm. We call the final outer radius $r_a$ the potential size, and the corresponding
47   radius of maximum winds $r_{\text{max}}$. The potential intensity $V_p$ and potential size $r_a$ are consistent with each
48   other as we used the potential intensity for both the CLE15 and W22 models, and the same environ-
49   mental variables such as the sea surface temperature $T_s$. The consistency of our modelling approach is
50   further enhanced by driving the ADCIRC model with the axisymmetric wind and pressure profile that
51   corresponds to the CLE15 output at this potential size $r_a$ and this potential intensity $V_p = V_{\text{gm}}$.

Figure 15 shows a single solution of the potential size calculation summarised in Figure 12. The solution marked as a cross is where the two models produce pressures at the radius of maximum wind-speed $p_{m1}$, $p_{m2}$ are within some threshold value $t$ of one another (taken arbitrarily as 1 Pa). These two curves are expected to cross because the energetic constraints of the W22 Carnot engine would reduce the central pressure deficit with higher $r_a$, and the dynamic constraints of the CLE15 radial profile would increase the central pressure deficit with higher $r_a$. We initially find the intersection by using the bisection method for simplicity, and because there was not an obvious way of calculating the gradient of pressure deficit by change in outer radius, $\frac{dp_{m2}}{dr_a}$, for the Chavas *et al.* (2015) radial profile.

To extend this further Figure 16 shows the curves calculated from the monthly average data from each August of a climate model ensemble member from 1850 to 2100, where 1850-2014 is from the historical simulation, and 2015-2100 is from the SSP-585 scenario. This is calculated given the conditions at the centre of the Gulf of Mexico. Both the W22 and CLE15 curves move in response the climate change and other factor so that their intersection also moves. We can see that over time the potential size increases, as the more recent years tend to be further to the right. There is a significant spread in the central pressure deficit $p_m$ where this solution is found. One unexplained problem introduced in Figure 2 was why the internal variability in the potential intensity $V_p$ was so much higher than the potential size $r_a$. It is perhaps possible that the change in both curves together somehow leads to a lower $r_a$ than you would expect given that we are assuming that $V_p$ can validly be used as one of the inputs to model to calculate $r_a$.

We use the inputs from the August monthly average because this around the peak of the hurricane season activity, and also around the peak for potential intensity, $V_p$. To illustrate this see Figure 13 where we show the variation of potential intensity and its inputs over ten years of a CMIP6 ensemble member for the point closest to New Orleans ($-90.25°$E, $29.25°$N). As shown, for this point the potential intensity, $V_p$, tends to peak in September, whereas the potential size, $r_a$, has flat peak from roughly June to September. Figure 14 shows the corresponding vertical profiles for some of the months from the first year of the ensemble member. As shown, the outflow level rises to roughly the temperature inversion pressure level in August 2015, from a much lower level in February 2015.
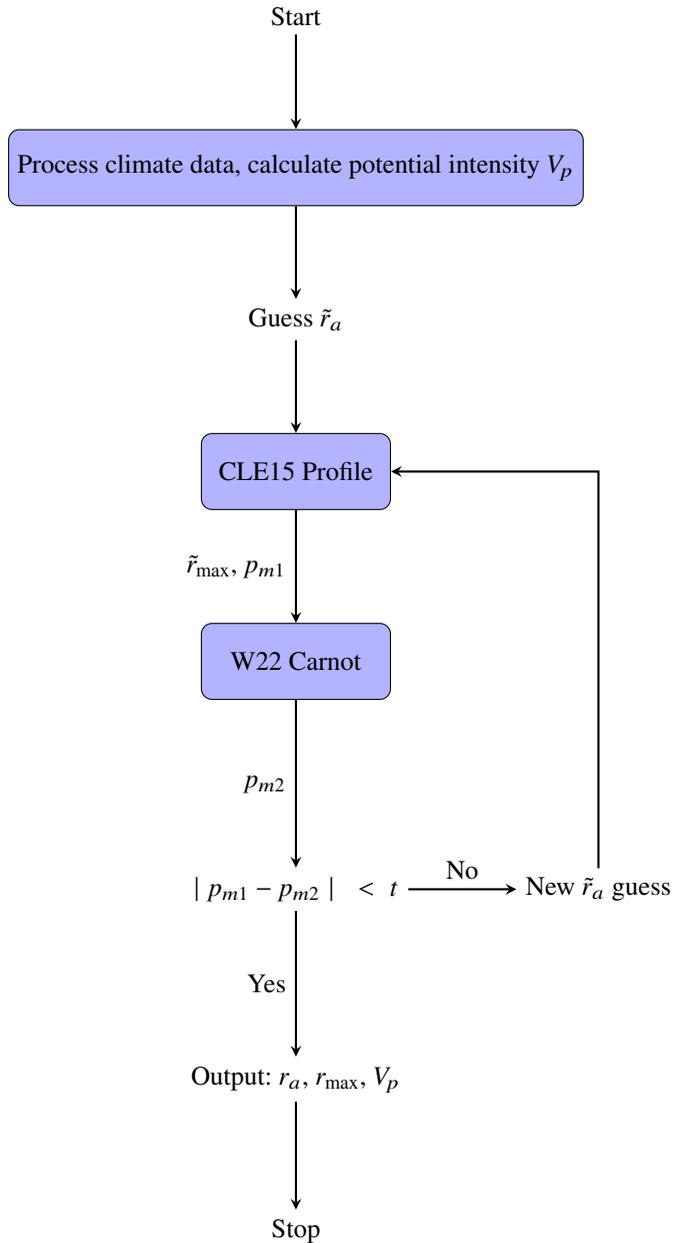
The calculation of potential size is quite slow, taking around 2 minutes per grid point per time point. This is partially because the CLE15 profile calculation code is written in `matlab` (Chavas 2022), and was ran by octave at each new guess of $\tilde{r}_a$, which involves a launching cost. To pass the input data to and from octave, a `json` file is saved and its name is passed. Each point can be parallelized onto a separate CPU, but this still makes the calculations too slow to easily run for large CMIP6 datasets. It should be possible to improve this in future by writing this purely in optimized python or another language.

## 7. ADCIRC model setup

We apply the atmospheric forcing in ADCIRC with netCDF4 files using the NWS=13 input setting. The inputs are the 10m zonal wind (U10) in m s$^{-1}$, the 10m meridional wind (V10) in m s$^{-1}$, and the surface pressure (PSFC) in mbar. The 10m windspeed $\left|\vec{U}_{10}\right|$ is calculated as a constant factor $V_{\text{reduc}} = 0.8$ of the gradient wind, $V(r)$, and the surface pressure is calculated as in Equation 9 using the gradient rather than the 10m wind. If each point in the mesh is at position $\vec{p}$ and the centre of the tropical cyclone is at position $\vec{p}_c$, we calculate the distance from the centre for each point, $r = ||\vec{p} - \vec{p}_c||$, and its angle $\phi = \arctan 2\left(p_x - p_{cx}, \ p_y - p_{cy}\right) - \pi/2$, so that the 10m wind vector becomes,

$$(U10(\vec{p}), \ V10(\vec{p})) = (\sin(\phi) \cdot V(r) \cdot V_{\text{reduc}}, \ \cos(\phi) \cdot V(r) \cdot V_{\text{reduc}}), \tag{11}$$

and defaults to $(0,0)$ if $r > r_a$. Similarly the pressure field is PSFC $= p(d)$, and defaults to $p_a$ when $r > r_a$.

**Figure 12.** *Flowchart for the calculation of the outer radius, or potential size, $r_a$, of a tropical cyclone that satisfies both the dynamical constraints of the Chavas* et al. *(2015) radial wind profile and the thermodynamic constraints of the Wang* et al. *(2022a) Carnot engine. The solution is for when the two models produce pressures at the radius of maximum windspeed $p_{m1}$, $p_{m2}$ are within some threshold value t of one another by varying the outer radius of the tropical cyclone $\tilde{r}_a$ until there is agreement, $|p_{m1} - p_{m2}| < t$, at the potential size $r_a$.*
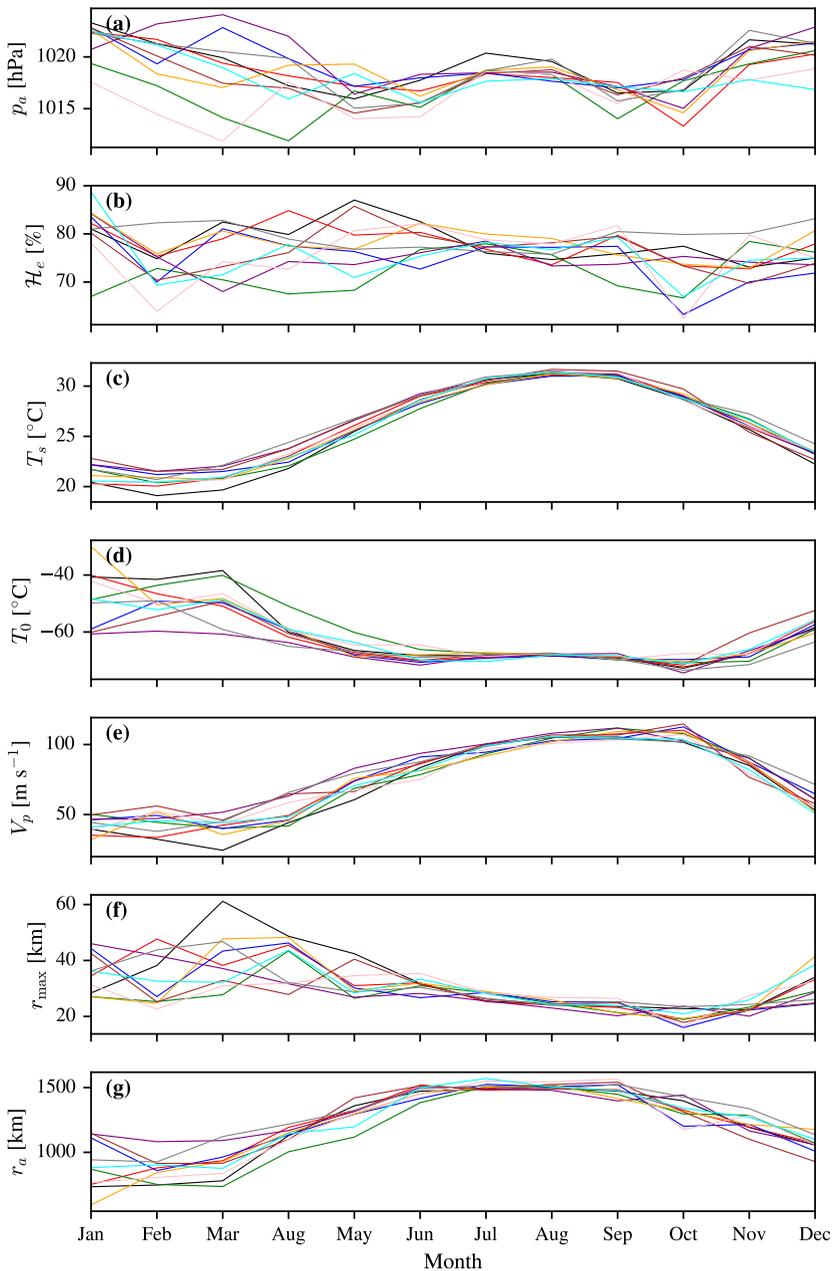
***Figure 13.*** *A comparison of the seasonal cycles for input and output variables for the ocean grid point centred closest to New Orleans in the CESM2-r4i1p1f1 SSP-585 CMIP6 ensemble member between 2015 and 2024 (inclusive). Each year is shown in a different color. The inputs include the sea surface temperature $T_s$, the mean sea level pressure $p_a$, the relative humidty $\mathcal{H}_e$, and the outflow temperature $T_0$. The outputs include the potential intensity at the gradient wind level $V_p$, the radius of maximum winds $r_{\max}$, and the potential size $r_a$.*
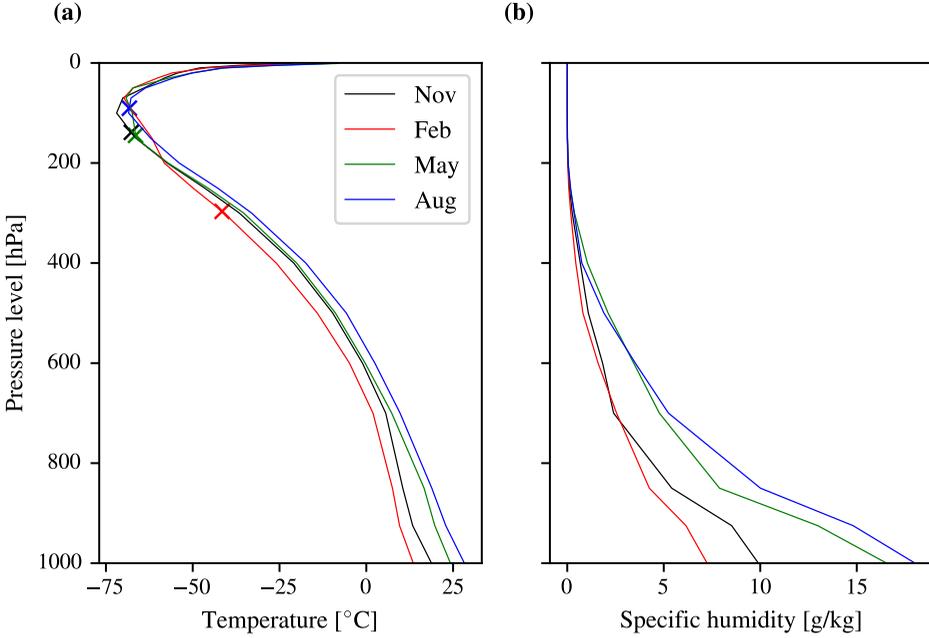
01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22



***Figure 14.*** *The vertical profiles of air temperature, t, and specific humidity, q, from a grid point point near New Orleans in February, May, August, and November 2015 SSP585-CESM2-r4i1p1f1. The crosses on the left panel are the outflow temperatures $T_0$ and the outflow levels OTL calculated by the tcPYPI package..*

The straight-line constant bearing trajectory of the tropical cyclone is calculated so that it makes hits the impact location $\vec{p}_i$ at the impact time $t_i$. The impact location $\vec{p}_i$ is defines as the observation location $\vec{p}_o$ minus the track displacement, $c$, so that $\vec{p}_i = (p_{ox} - c, p_{oy})$. The TC trajectory passes at translation speed $V_t$ at track bearing $\chi$ from due north, so that the position of the centre of the track, $p_c$, at time $t$ is $\vec{p}_c = \vec{p}_i + (t - t_i) \cdot V_t \cdot (\sin(\chi), \cos(\chi))$.

These are input on two rectangular lat-lon grids: a coarser static grid ($\frac{1}{8}^{\circ}$), and a higher resolution moving grid centred on the tropical cyclone centre $\frac{1}{80}^{\circ}$. The purpose of the higher resolution moving grid is to better resolve the centre of the tropical cyclone and decrease the errors brought about by inter-polation in each time slice, as the interpolation algorithm used between input timesteps is movement aware. Using the NWS=13 setup also makes the ADCIRC run more flexible so that any cyclone, includ-ing arbitrarily asymmetric cyclones could be applied, alongside a background climatology, so that we can easily extend this framework to more realistic conditions.

Wang *et al.* (2022b) found that forcing the ADCIRC model with the axisymmetric CLE15 profile, with a background wind model for the asymmetric component (Lin & Chavas (2012), referred to as LC12), significantly outperformed more standard parametric profiles such as the revised Holland profile (Holland *et al.* 2010) in the same setup, particularly for the more extreme storm surges. Chaigneau *et al.* (2024) showed that forcing ADCIRC with the CLE15 profile outperformed forcing it with ERA5 data, and was comparable to parametric models such as the generalized asymmetric Holland model (GAHM).

Each of our runs involves 7 virtual days of simulation time. The tides have been turned off, and there is one day of linear ramping at the start of the meteorological data to avoid a hard shock of the new imposed conditions from creating transient effects. The run is parallelized across a node with 128 CPUs and each ADCIRC run takes around 3 minutes.
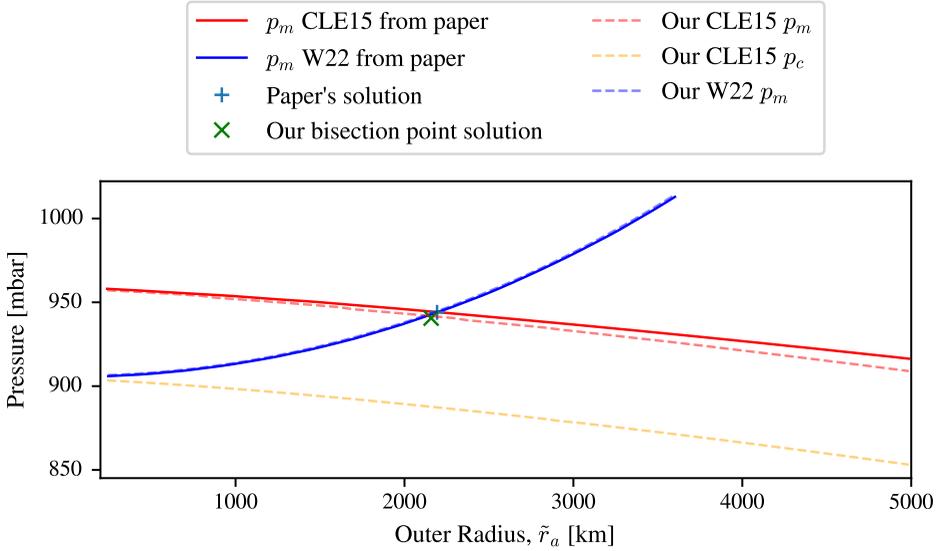
*Figure 15. Finding a solution that satisfies Wang* et al. *(2023)'s Carnot engine (W22) and Chavas et al. (2015)'s profile wind profile (CLE15). I compare the curves produced by varying the outer radius from our implementation to those extracted from Wang* et al. *(2022a)'s Figure 4a using the WebPlot-Digitizer software. The cross (×) marks the intersection found using bisection as outlined in the text, whereas the plus (+) marks the solution reported in their paper. Assuming the CLE15 wind profile to be cyclogeostropically balanced means that the larger the outer radius $r_a$ the larger the central pressure deficit. Assuming that the W22 Carnot cycle can only extract so much energy means that the larger $r_a$ the smaller the central pressure deficit. Therefore, we expect one $r_a$ that satisfies both constraints. This graph was generated with $V_{sg} = 83/\gamma_{sg}$ m s$^{-1}$, $\gamma_{sg} = 1.2$, $\rho_a = 1.225$ kg m$^{-3}$, $\mathcal{H}_e = 0.9$, $p_a = 1015$ mbar, $T_0 = 200$ K, $T_s = 299$ K, $L_v = 2,500,000$ J kg$^{-1}$, $w_{cool} = 0.002$ m s$^{-1}$, $C_d = 0.0015$, $\frac{C_k}{C_d} = 1$. We assumed the pressure could be integrated inwards with the density decreasing as an ideal gas in isothermal conditions (at $T_s$) using a trapeziodal routine in* `scipy.integrate`*. The pressure from the CLE15 profile at the maximum winds is slighly lower in our solution than theirs, which may be caused by our choice of integration method, or that they used a higher (unreported) density. $p_c$ is the central pressure of the tropical cyclone in the CLE15 model which is roughly a constant 50 mbar lower than the pressure at the radius of maximum winds, $p_m$.*

For the Bayesian optimization experiments in the New Orleans region we apply the constraints that the tropical cyclone broadly comes from the south, $\chi \in (-80°, \ 80°)$, the tropical cyclone hits the region of the observation point, $c \in (-2°\text{E}, \ 2°\text{E})$, and the translation speed is positive, and up to twice the translation speed of Hurricane Katrina at Landfall, $V_t \in (0 \text{ m s}^{-1}, \ 15 \text{ m s}^{-1})$. Each of these decisions is somewhat arbitrary, particularly for the upper bound translation speed, $V_t$, which can become much faster. Translation speed could also interact in a complex way with both the size and intensity of a tropical cyclone. If 15 m s$^{-1}$ is too low an upper bound, this would likely lead to an underestimate in the potential height for the open coastlines, but these are not the areas of the coast that were at the risk of the highest surges in Section 3.
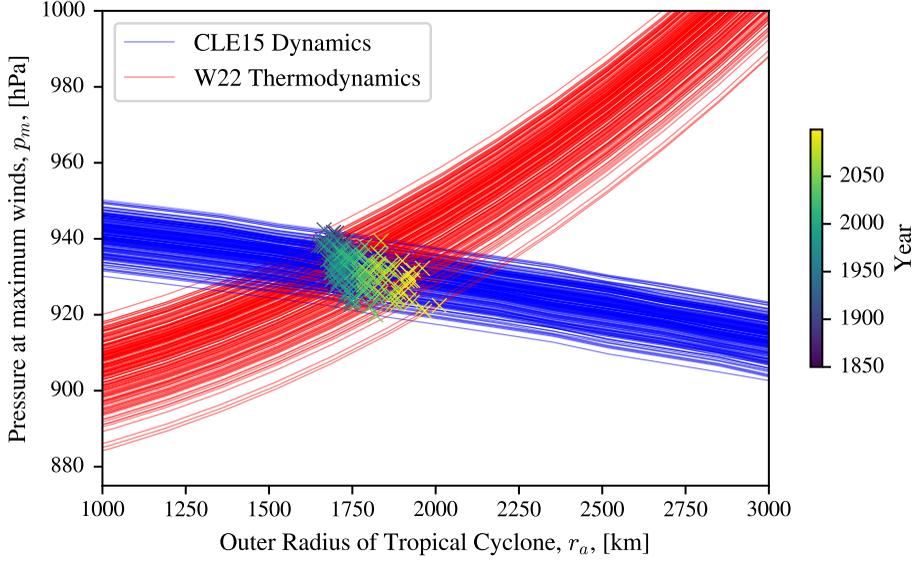
***Figure 16.*** *Finding solutions that satisfies Wang et al. (2022a)'s Carnot engine (W22) and Chavas et al. (2015)'s profile (CLE15). Each set of lines corresponds to the solutions found for different mean monthly properties in a different year. The crosses are the intersections between the different corresponding years, and these crosses are colored by the year of the pair of curves.*

## 8. Investigating GEV fit sensitivity to the scale and shape parameters

In Figure 17, we compare the sizes of the bias and 5%–95% resampling ranges from the max known fit (I) and the max unknown fit (II) for the same return levels highlighted in Figure 10. We vary the values of the shape parameter, $\gamma$, and the scale parameter, $\beta$ chosen for the true GEV systematically for a plausible number of observations, $N_s = 50$. We keep the resamples constant $N_r = 600$, and vary $\gamma$ between -0.1 and -0.4, and $\beta$ between 0.2 m and 1.5 m. The biases in the estimate of the return value is much smaller for (I) than (II) (Figure 17(a),(f),(c),(h)), and the 5%–95% resampling ranges are also smaller (Figure 17(b),(g),(d),(i)). This is shown systematically in Figure 17(e)&(j) where the natural logarithm of the 5%–95% range for (II) divided by (I) is almost always positive. This illustrates that for a plausible size of the observational dataset, assuming block maxima, the 1 in 500 year return value in particular can be estimated much more accurately if the upper bound were known ahead of time, independent of the values of the shape and scale parameters, $\gamma$ and $\beta$.
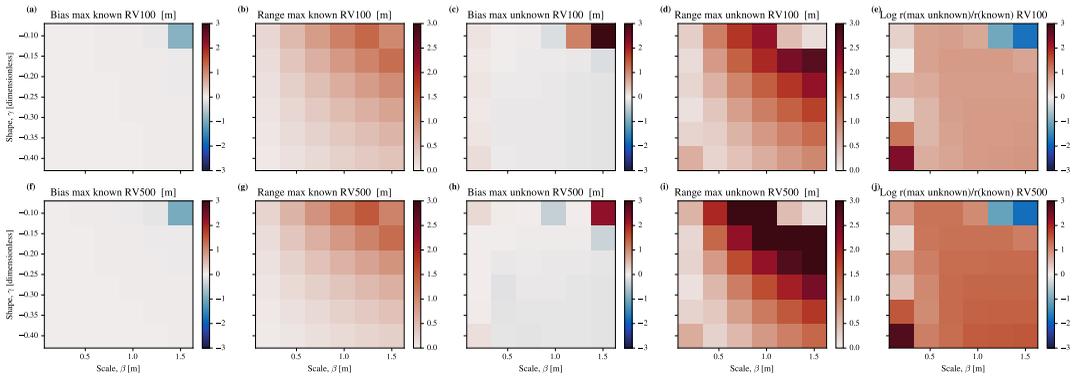
**Figure 17.** *We also experiment with varying the shape parameter γ and the scale parameter β whilst keeping the maximum height at $z^* = 7$ m. The number of samples is $N_s = 50$, and the number of resamples is $N_r = 600$. Not knowing the upper bound leads to larger biases and larger ranges, even more true for RV500 than RV100. Panel (a) and (f) shows the bias in the estimate of 1 in 100 and 1 in 500 year return value respectively compared to the true value. Panels (b) and (g) show the range between the 5% and 95% estimate of those return levels, estimated by resampling $N_r = 600$ times. Panels (c) and (h), (d) and (i) show the same metrics but for the case where the maxima was not known ahead of time, where the biases and ranges are generally larger than where the max is known. To demonstrate this effect quantitatively for the ranges, we plot the natural logarithm of the range calculated for the max unknown fits, divided by the range calculated by the max known fits in panels (e) and (j). This value is positive (red) where the range is reduced for the max known fit compared to the max unknown fit. For almost all of the parameter space knowing the upper bound seems to decrease the range, apart from for β = 1.5 m and γ = −0.1, but the bias is substantially larger there instead.*