Supplementary information: Extreme Atlantic hurricane seasons made more likely by ocean warming

Peter Pfleiderer[a,b,*], Shruti Nath[a,c], and Carl-Friedrich Schleussner[a,b]

[a]Climate Analytics, Berlin, Germany
[b]Humboldt University, Geographie, IRI-Thesys, Berlin, Germany
[c]Eidgenössische Technische Hochschule Zürich, Zürich, Switzerland
[*]Corresponding author: peter.pfleiderer@climateanalytics.org

October 5, 2021

Tropical Cyclone Season Emulator

0.1 Motivation

The aim of this emulator is to generate plausible Atlantic hurricane seasons for given large-scale atmospheric conditions - characterized by daily weather patterns - and sea surface temperatures (SST) averaged over the main development region (MDR). With this simple separation into dynamic (weather patterns) and thermodynamic (SSTs) input variables we can analyse simple counterfactual scenarios with altered SSTs to estimate the influence of warming oceans on hurricane activity.
Note that weather patterns are to some extend influenced by SSTs in the region which makes it impossible to fully disentangle dynamic and thermodynamic contributions to hurricane activity. Here we take the lack of any robust changes in atmospheric circulation in the region as basis for the assumption that similar sequences of weather patterns would occur in a preindustrial climate. With this assumption we study the question to which extend the observed warming of the tropical Atlantic Ocean has contributed to an increase in hurricane activity.

0.2 Emulator Design

We construct a probabilistic emulator that creates a series of storms with maximum sustained wind speeds for each day. Hurricanes are rare events and their formation and intensification involves complex physical processes. In our emulator we break these processes down into three components that are fully independent from each other:

1. storm formation
2. storm duration
3. daily storm intensity

In these components the daily weather pattern and the daily SST average over the MDR (only for the daily storm intensity component) slightly alters the probabilities for a new storm formation or for the intensification of an existing storm. The design of the emulator is schematically shown in figure S11.
The emulator does not simulate storm locations which is a strong limitation as the effect of a certain weather pattern on a given storm might strongly vary depending on the storm location. This limitation is however necessary as there is not enough data to evaluate the effect of weather patterns on storms everywhere in the Atlantic basin. Due to this limitation there remains a strong stochastic element in our emulator which reflects the information that would be required for the simulations on top of the daily weather patterns and regionally averaged SSTs.

1 Emulator Components

For each of the emulator components heuristic methodological decision had to be taken. In the following we explain the underlying assumptions and show how the emulator performs compared to alternative methodological choices.

1.1 Storm Genesis

Storm formations are rare (roughly XX per season) but they predominantly occur during favorable weather conditions including low vertical wind shear, high relative humidity an the lower troposphere and the existence of some kind of perturbation. Tropical storm formations also require local SSTs of at least 26C.

Assumptions:

1. Weather patterns can favour or hamper storm formations
   (a) Persistent weather conditions can further increase or decrease formation probabilities
2. regionally averaged SSTs contain only limited information on storm formation probabilities

Based on these assumptions we estimate the probability of a storm formation event P_{gen} on a day d with weather pattern $w(d)$ as:

$$P_{gen}(d) = P_{obs}(gen|w(d)) * \frac{\sqrt{P_{obs}(gen_{next\ day}|w(d-1)) * P_{obs}(gen_{2\ days\ after}|w(d-2))}}{P_{obs}(gen|all)} \tag{1}$$

The first factor $P_{obs}(gen|w(d))$ is the observed probability of storm formations for the given weather pattern. The second factor includes the probabilities of storm formations one and two days after the weather pattern that occurred one and two days earlier respectively. These probabilities receive less weight and the factor is normalized by a dividing by the overall observed storm formation probability.

1.1.1 Tested variations

vG0 – Simplest approach: The probability for a new storm is the number of genesis events during the weather pattern of the day w divided by the number of observed days with that weather pattern:

$$P_{gen}(d) = P_{obs}(gen|w(d)) \tag{2}$$

vG1 - giving more weight to the weather history: As in the main component, but multiplied by the probability of finding a genesis event on the following day of the weather pattern of the day before and respectively two days before. These additional probability factors are normalised by dividing by the overall observed genesis probability squared:

$$P_{gen}(d) = P_{obs}(gen|w(d)) * \frac{P_{obs}(gen_{next\ day}|w(d-1)) * P_{obs}(gen_{2\ days\ after}|w(d-2))}{P_{obs}(gen|all)^2} \tag{3}$$

vG2 - nearest neighbours with weather and SST: For each combination of SST and weather the 100 nearest neighbours in the observations are taken and the genesis probability is estimated from these observations.

We use the Euclidean distance metric and standardise our variables for the distance calculation.

$$D(d_i, d_j)^2 = \frac{(SST(d_i) - SST(d_j))^2}{\sqrt{\frac{1}{N}\sum_m (SST(d_m) - \overline{SST})^2}} + \frac{(w(d_i) - w(d_j))^2}{\sqrt{\frac{1}{N}\sum_m (w(d_m) - \overline{w})^2}} \tag{4}$$

For the weather patterns, which are not a continuous variable, we consider their coordinates in the SOM grid as locations and calculate differences between weather patterns as the sum of the squared differences in row and column numbers.

$$w(d_i) - w(d_j) = \sqrt{(w_{row}(d_i) - w_{row}(d_j))^2 + (w_{col}(d_i) - w_{col}(d_j))^2} \tag{5}$$

1.1.2 Conclusions

The genesis component reproduces the inter-annual variations in storm formations adequately. Removing information about the preceding weather patterns slightly reduces the skill of the emulator. Adding regionally averaged SSTs does however lead to considerably worse representation of inter-annual variability.

1.2 Storm duration

There are numerous processes that can weaken and eventually dissipate tropical cyclones. The most common end of a TC is landfall. As we don't have information about the location of storms in our emulator, estimating the duration of a storm is challenging.

Assumptions:

1. The most relevant end of storms is landfall
2. The time a storm has before making landfall depends on its formation location
3. The formation location is to some extend influenced by the weather patterns

As shown in figure S5, storms predominantly form in the eastern part of the main development region (MDR) for weather patterns 6, 7, 10, 11 and 15. Under favorable conditions, storms forming in the eastern part of the MDR have a long way before encountering land and thereby a potential for lasting longer than storms forming in the western part of the MDR. We incorporate this (assumed) dependence of storm duration on weather patterns on the day of storm formation by sampling the duration D of a storm s from a gaussian kernel estimate f_g of all storms that have formed during the weather pattern during which the storm has formed $w(d_f)$ and all neighboring weather patterns.

$$D(s) = f_g(D_{obs}[w_{row} - w_{row}(d_f) < 2 \ \& \ w_{col} - w_{col}(d_f) < 2]) \tag{6}$$

Figure S6 shows the observed duration of storms that formed during a weather pattern and the distributions from which storm durations are sampled in the emulator. The assumption, that storm durations vary depending on the weather pattern during which a storm forms holds and is adequately reproduced by the storm duration component.

1.2.1 Tested variations

vD0 – independent of weather patterns: The duration of a storm is sampled from the probability distribution of all observed storms:

$$D(s) = f_g(D_{obs}) \tag{7}$$

vD1 – without averaging over neighboring weather patterns: The length of a storm is sampled from the probability function of all storms that emerged during the same weather pattern:

$$D(s) = f_g(D_{obs}[w = w(d_f)]) \tag{8}$$

1.2.2 Conclusions

The effect of weather conditions on individual storm durations is limited (see below). Based on our assumptions we are able to estimate plausible durations for the simulated storms. The tested variations of the component have limited influence on the emulator performance.
(Note that systematically different approaches for the estimation of storm durations are possible. For example one could estimate the probability for the end of a storm on each day based on surrounding conditions. In such an approach the lack of storm locations would be a critical flaw.)

1.3 Storm intensity

Storm intensity is quantified through the daily maximum sustained wind speed of the storm. We use the following assumptions for our daily storm intensity emulations:

1. Intensification can be favoured or hampered by specific weather patterns
2. Warmer SSTs in the main development region generally favour TC intensification
   (a) The relationship between SSTs and storm strength con be represented by a quantile regression
3. The strength of TCs is auto-correlated and depends on the storm strength of the day before

Assessing probability density functions for daily max. wind speeds intensities for all possible combinations of weather patterns, SSTs, and storm histories is challenging given the insufficient number of storm observations. Therefore, instead of estimating a PDF for the daily max. intensity from all observations that match to certain conditions (e.g. weather pattern 6, 28°C SST and 60kts wind speed on the day before) we estimate the wind speed PDF from the 100 storm observations that are most similar to these conditions.
Furthermore, the range of observed SSTs is limited with only few observations in the range of pre-industrial SSTs. We therefore assume a linear relationship between regionally averaged SSTs and storm intensities: for a given SST_{target} we transform all observed storm intensities to artificial pseudo-intensities $v_{shifted}$ using the slope β_τ of the next quantile τ below the observed storm strength $v(s, d)$:

$$\begin{aligned} &v_{shifted}(s, d, SST_{target}) = v(s, d) + \beta_{\tau(s,d)} * (SST_{target} - SST_{obs}(d)) \\ &\tau(s, d) = min(\tau : v(s, d) > \beta_\tau + c_\tau) \end{aligned} \tag{9}$$

As a distance metric for the 100 nearest neighbors and the remaining variables (weather pattern and storm strength on the day before) we use an Euclidean distance with standardized variable:

$$\begin{aligned} D(d_i, d_j)^2 = &\frac{(w(d_i) - w(d_j))^2}{\sqrt{\frac{1}{N}\sum_m (w(d_m) - \overline{w})^2}} \\ &+ \frac{(v_{shifted}(d_i - 1) - v_{shifted}(d_j - 1))^2}{\sqrt{\frac{1}{N}\sum_m (v_{shifted}(d_m) - \overline{v_{shifted}})}} \end{aligned} \tag{10}$$

For the weather pattern, which are not a continuous variable, we consider their coordinates in the SOM grid as locations and calculate differences between weather patterns as the sum of the squared differences in row and column numbers.

$$w(d_i) - w(d_j) = \sqrt{(w_{row}(d_i) - w_{row}(d_j))^2 + (w_{col}(d_i) - w_{col}(d_j))^2} \tag{11}$$

1.3.1 Tested versions

vI0 – 100 nearest neighbours Straight forward 100 nearest neighbors with conditions being characterized by weather patterns, SSTs and storm strengths on the day before. This is a simpler variation that works without the quantile regression between storm strengths and SSTs.

$$\begin{aligned} D(d_i, d_j)^2 = & \frac{(SST(d_i) - SST(d_j))^2}{\sqrt{\frac{1}{N}\sum_m (SST(d_m) - \overline{SST})^2}} \quad + \frac{(w(d_i) - w(d_j))^2}{\sqrt{\frac{1}{N}\sum_m (w(d_m) - \overline{w})^2}} \\ & + \frac{(v(d_i - 1) - v(d_j - 1))^2}{\sqrt{\frac{1}{N}\sum_m (v(d_m) - \overline{v})}} \end{aligned} \tag{12}$$

vI1 – 20 nearest neighbors 20 nearest neighbors based on weather patterns and storm intensities on the day before and SSTs.

vI1 – no sst dependence 100 nearest neighbors based on weather patterns and storm intensities on the day before.

vI3 - no weather dependence As the main component but without the dependence on weather patterns.

vI4 - no storm history Intensities are estimate irrespective of the intensity of the storm on the day before.

1.3.2 Conclusions

The storm intensity component is the most complex component of the emulator. With our nearest neighbor approach with the chosen variables and the quantile regression between SSTs and storm intensities we are able to reproduce the inter-annual variability of ACE hurricane days and major hurricane days.
The initial idea for this component of the emulator was to use a nearest neighbors approach (vI0). As shown in figure S12 and S13 there is simply not enough data to find close enough neighbors for all possible combinations of weather pattern, SST and intensities on the day before. This problem is most pronounced for strong storms under unfavorable weather conditions. Using 20 nearest neighbors (vI1) instead of 100 only slightly improves this problem.

Without any SST dependence (vI2), the simulations have a considerably larger trend in residuals of hurricane numbers, major hurricane numbers and ACE (see table S3 and figure S22-S24).

Omitting the information about the daily weather conditions (vI3) slightly reduces the representation of inter-annual variability (see table S2).

Removing the intensity of storms on the day before (vI4) from the nearest neighbor classification considerably reduces the representation of daily storm evolution (see figure S26-S28).

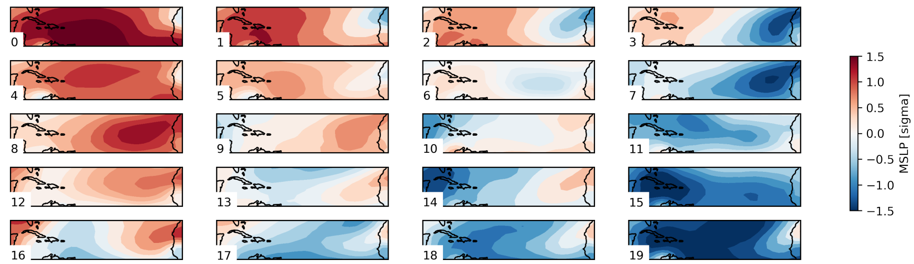

Figure S1: Mean sea level pressure (MSLP) anomalies for the 20 weather patterns expressed in standard deviations.

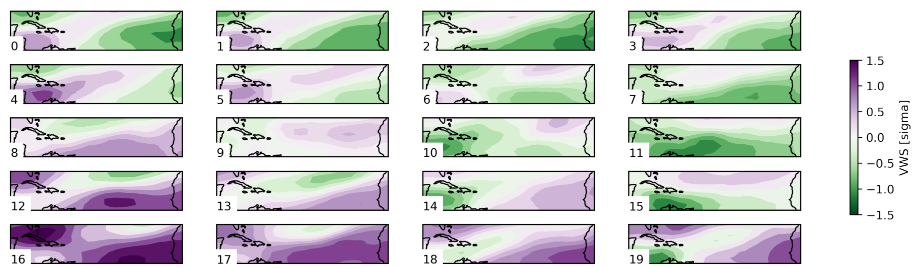

Figure S2: Vertical wind shear (VWS) anomalies for the 20 weather patterns expressed in standard deviations.

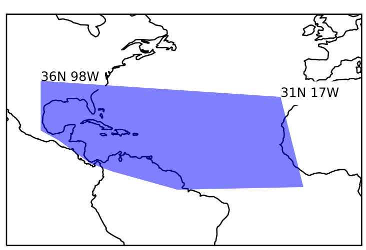

Figure S3: Storm observations within this polygon are considered in the analysis.

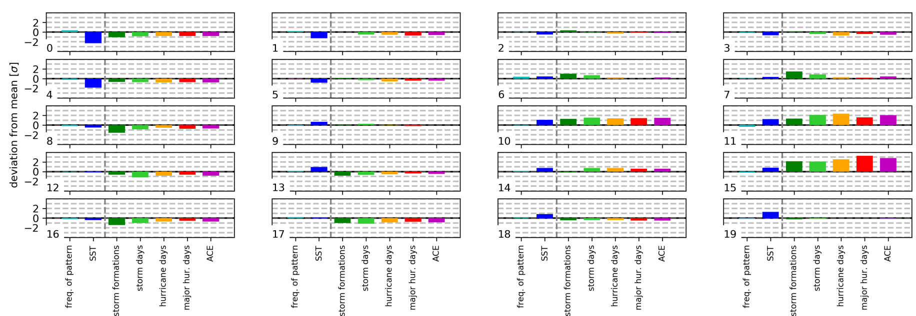

Figure S4: Tropical storm statistics for the 20 weather patterns expressed as relative deviations from the average.

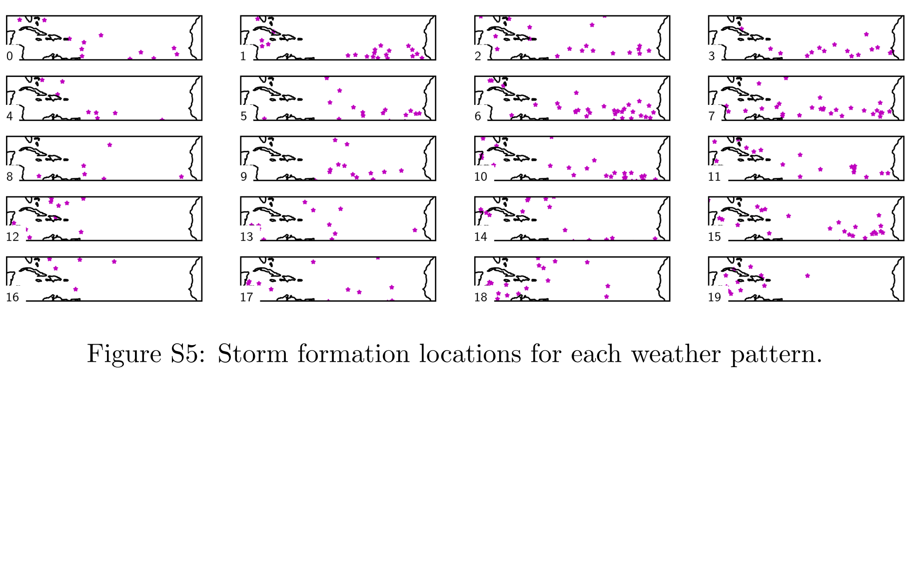

Figure S5: Storm formation locations for each weather pattern.

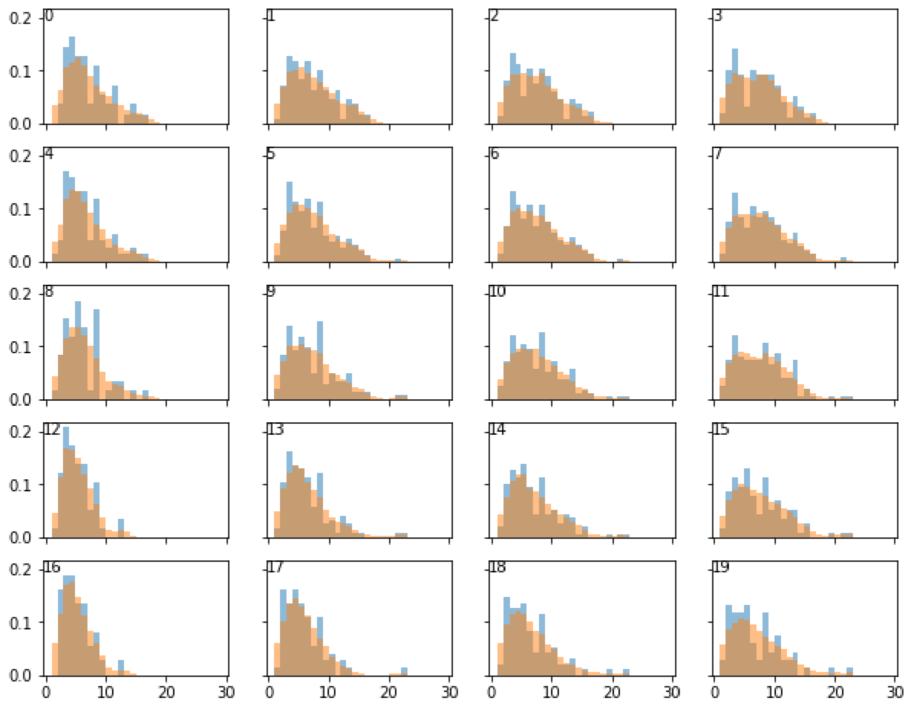

Figure S6: Histograms of observed storm durations for each weather pattern (blue). The orange histograms show the distributions from which storm durations are sampled in the emulator.

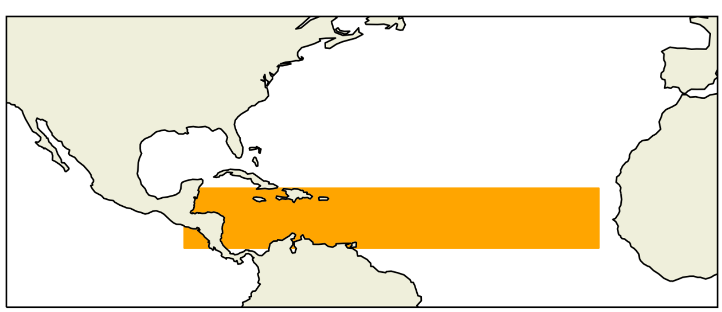

Figure S7: Atlantic main development region (MDR).

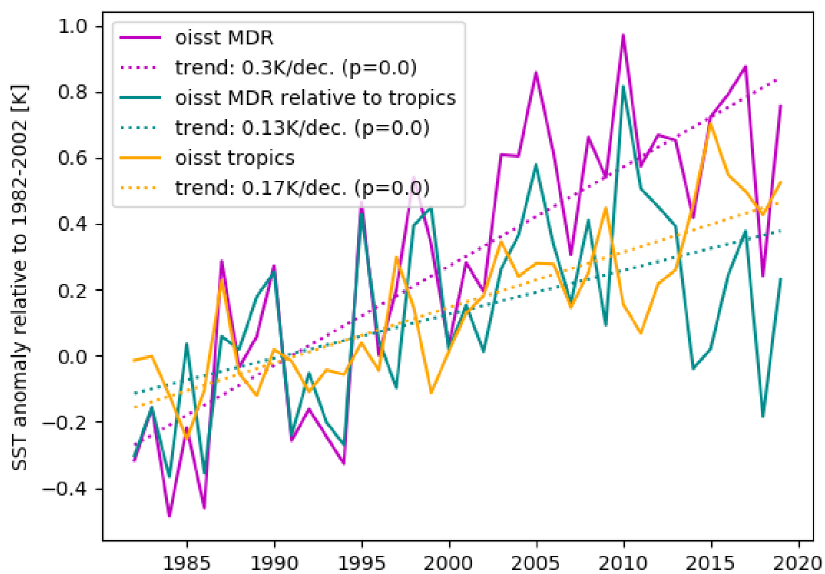

Figure S8: SST anomalies to the 1982-2002 average from OISST seasonally averaged over August-October for the MDR (pink), the tropics 30S-30N (orange) and MDR relative to the tropics (turquoise). Linear trends are indicated in the legend.

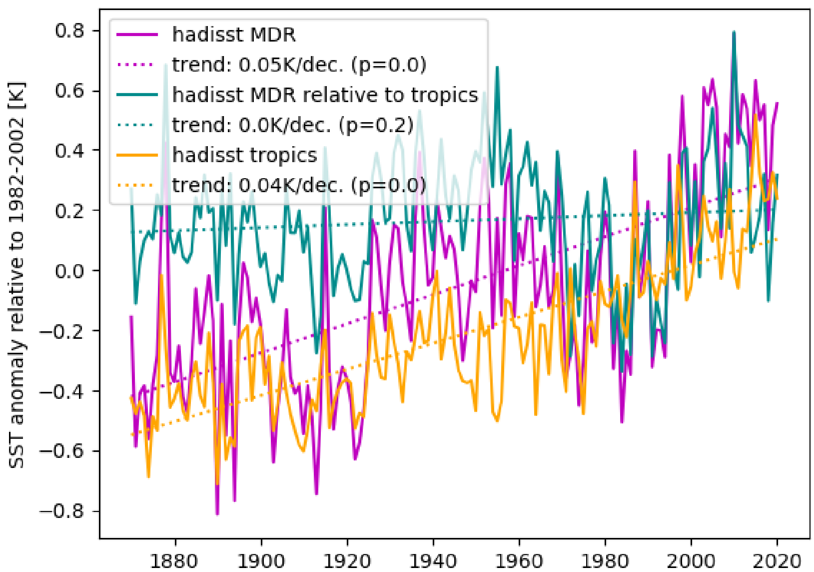

Figure S9: SST anomalies to the 1982-2002 average from HADISST seasonally averaged over August-October for the MDR (pink), the tropics 30S-30N (orange) and MDR relative to the tropics (turquoise). Linear trends are indicated in the legend.

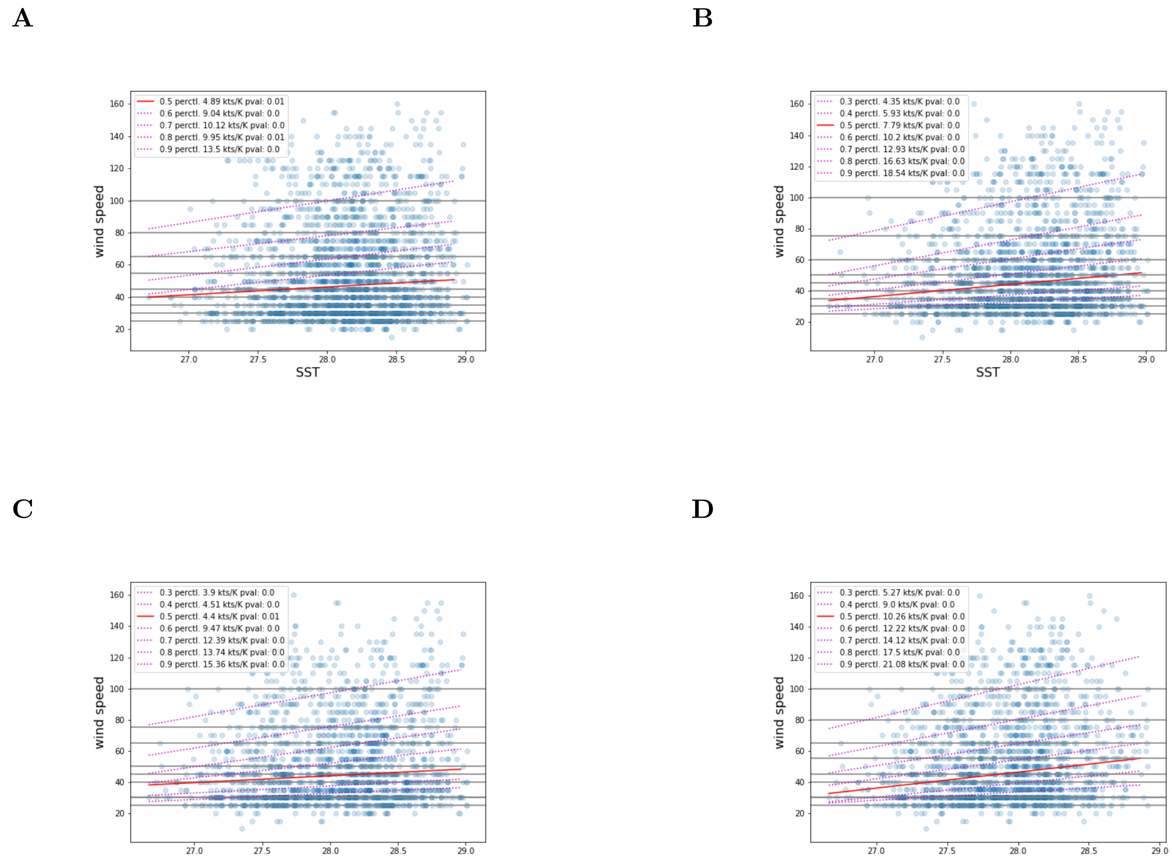

Figure S10: Quantile regression between storm intensities and SSTs over the MDR. Slopes for each quantile and the p-values of the regression are indicated in the legend for quantiles for which the regression was significant at the 90% level. A: Years in the period 1989-2018. B: Years in the period 1982-1988 and 1999-2018. C: Years in the period 1982-1998 and 2009-2018. D: Years in the period 1982-2008

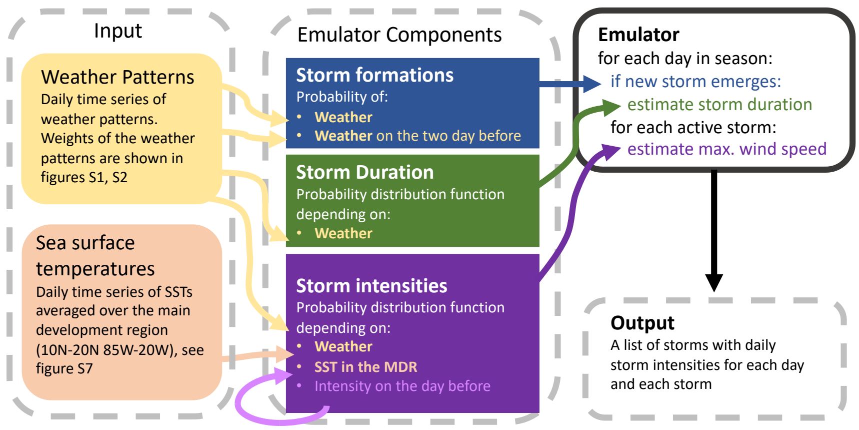

Figure S11: Schematic of the emulator design.

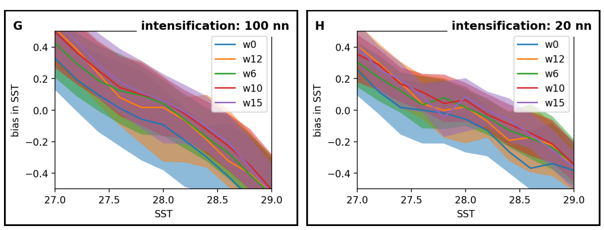

Figure S12: Deviation in SST from the desired SST for different sizes of nearest neighbors.

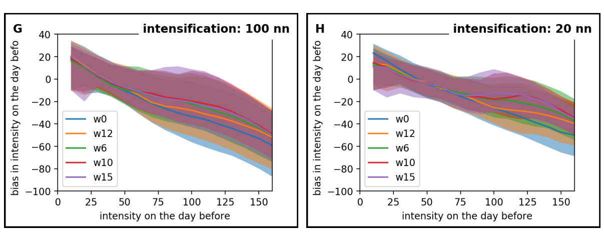

Figure S13: Deviations from the desired storm intensity on the day before for different sizes of nearest neighbors.

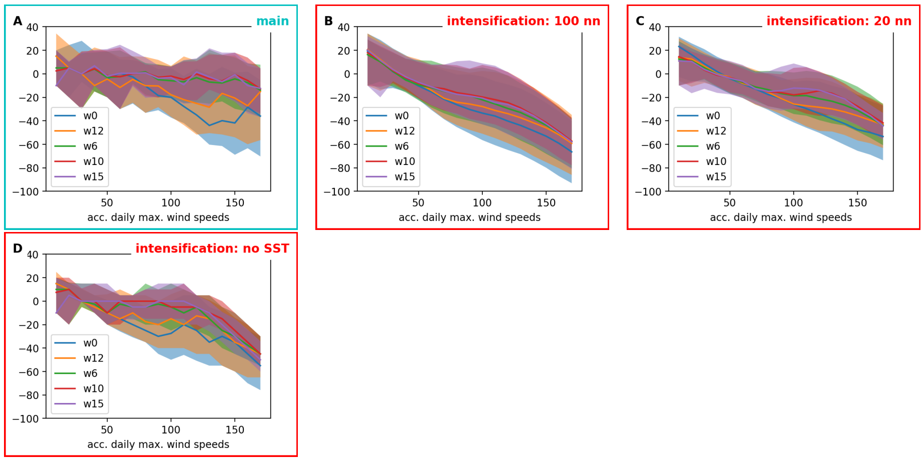

Figure S14: Deviations from the desired storm intensity on the day before for the 100 nearest neighbors.

Table S1: Root mean squared deviation for different indicators and versions of the emulator

version	storm formations	storm days	hurricanes	major hurricanes	ACE
main	3.08	21.57	1.67	1.14	38.34
formation: no lag	3.17	23.01	2.07	1.33	46.51
formation: equal weight	3.08	21.46	1.72	1.14	38.54
formation: NN weather + SST	3.46	24.2	2.17	1.35	47.66
duration: random	3.09	21.67	1.86	1.22	41.83
duration: no neighbors	3.1	21.27	1.68	1.09	36.83
intensification: 100 nn	3.05	21.55	1.77	1.22	41.84
intensification: 20 nn	3.09	21.47	1.94	1.24	40.11
intensification: no SST	3.08	21.26	1.85	1.23	40.94
intensification: no weather	3.16	21.58	2.19	1.21	42.14
intensification: no history	3.06	21.46	2.96	2.63	37.18

Table S2: Pearson correlation coefficients between observations and the median of 1000 simulations for different indicators and versions of the emulator.

version	storm formations	storm days	hurricanes	major hurricanes	ACE
main	0.61	0.68	0.8	0.74	0.77
formation: no lag	0.59	0.72	0.75	0.68	0.76
formation: equal weight	0.61	0.68	0.78	0.73	0.77
formation: NN weather + SST	0.41	0.61	0.7	0.69	0.69
duration: random	0.6	0.67	0.76	0.74	0.77
duration: no neighbors	0.6	0.69	0.79	0.76	0.78
intensification: 100 nn	0.61	0.68	0.78	0.67	0.75
intensification: 20 nn	0.6	0.68	0.74	0.65	0.76
intensification: no SST	0.61	0.68	0.76	0.68	0.73
intensification: no weather	0.59	0.67	0.74	0.72	0.72
intensification: no history	0.61	0.68	0.81	0.77	0.76

Table S3: Linear trends in residuals between observations and the median of 1000 simulations for different indicators and versions of the emulator.

version	storm formations	storm days	hurricanes	major hurricanes	ACE
main	-0.0	-0.64*	-0.02	-0.01	-0.53*
formation: no lag	-0.03	-0.83*	-0.03	-0.03	-0.71*
formation: equal weight	-0.0	-0.67*	-0.02	-0.02	-0.54*
formation: NN weather + SST	0.0	-0.75*	-0.04	-0.03	-0.74*
duration: random	-0.0	-0.75*	-0.03	-0.02	-0.61*
duration: no neighbors	0.0	-0.58*	-0.02	-0.01	-0.5*
intensification: 100 nn	0.0	-0.66*	-0.05*	-0.03*	-0.62*
intensification: 20 nn	0.0	-0.64*	-0.06*	-0.04*	-0.61*
intensification: no SST	0.0	-0.65*	-0.06*	-0.04*	-0.6*
intensification: no weather	0.01	-0.66*	-0.05	-0.03	-0.6*
intensification: no history	0.0	-0.66*	-0.01	0.02	-0.46*

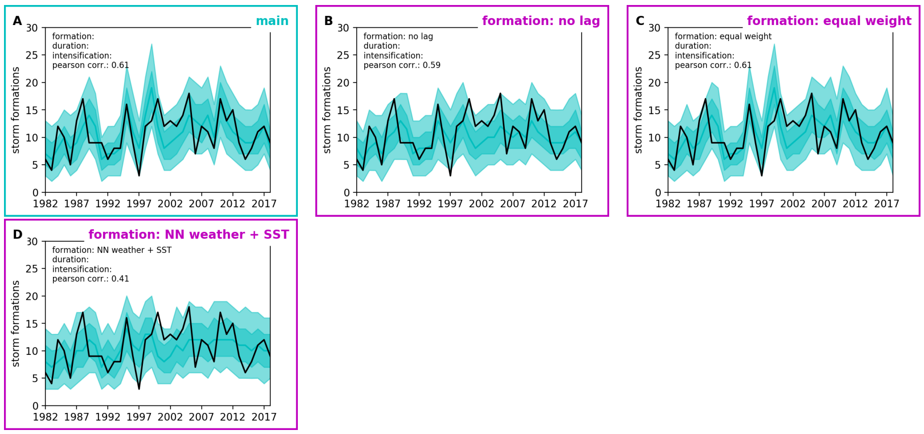

Figure S15: Simulations of seasonal storm formations for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

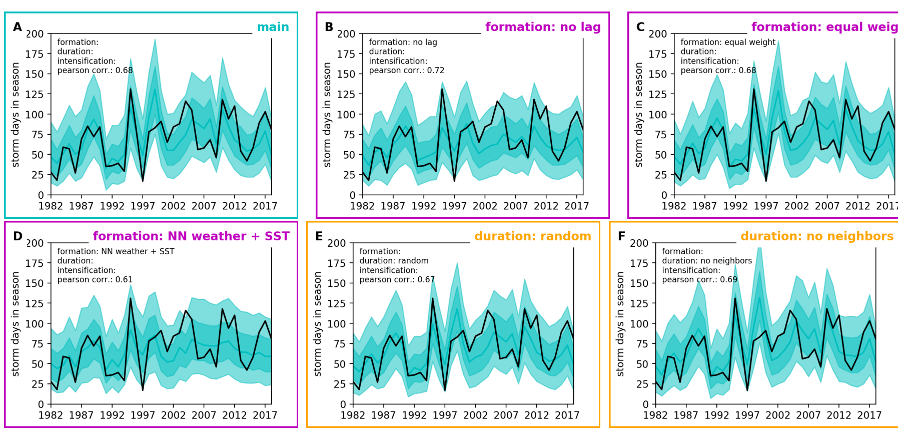

Figure S16: Simulations of seasonal storm days for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

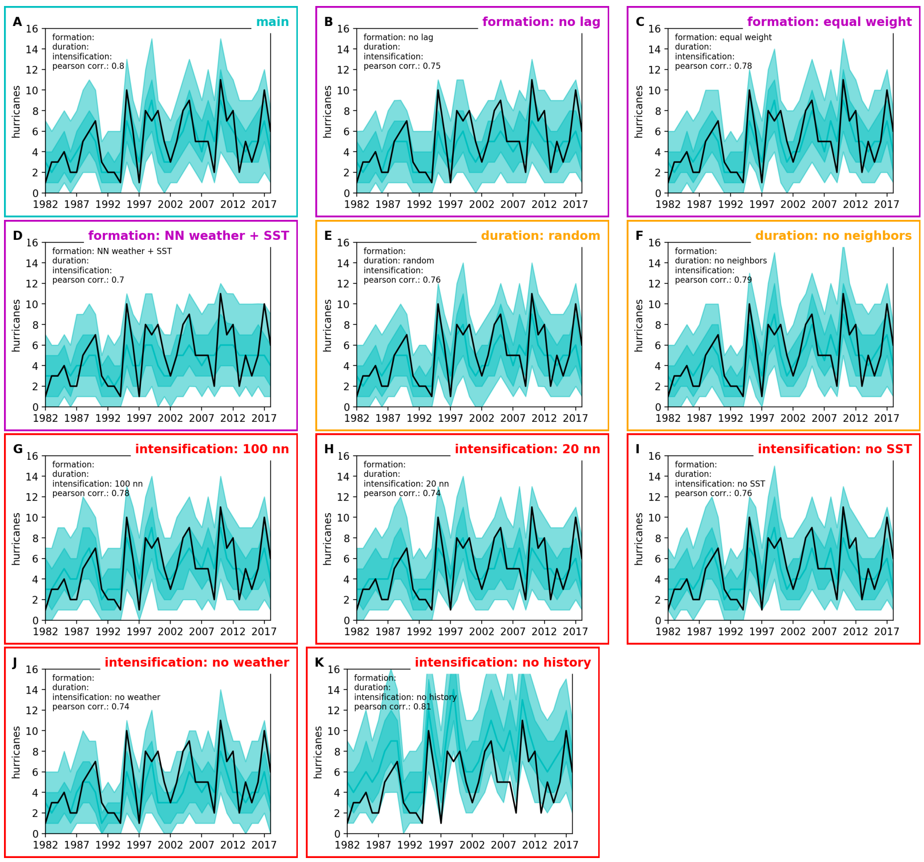

Figure S17: Simulations of seasonal hurricanes for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

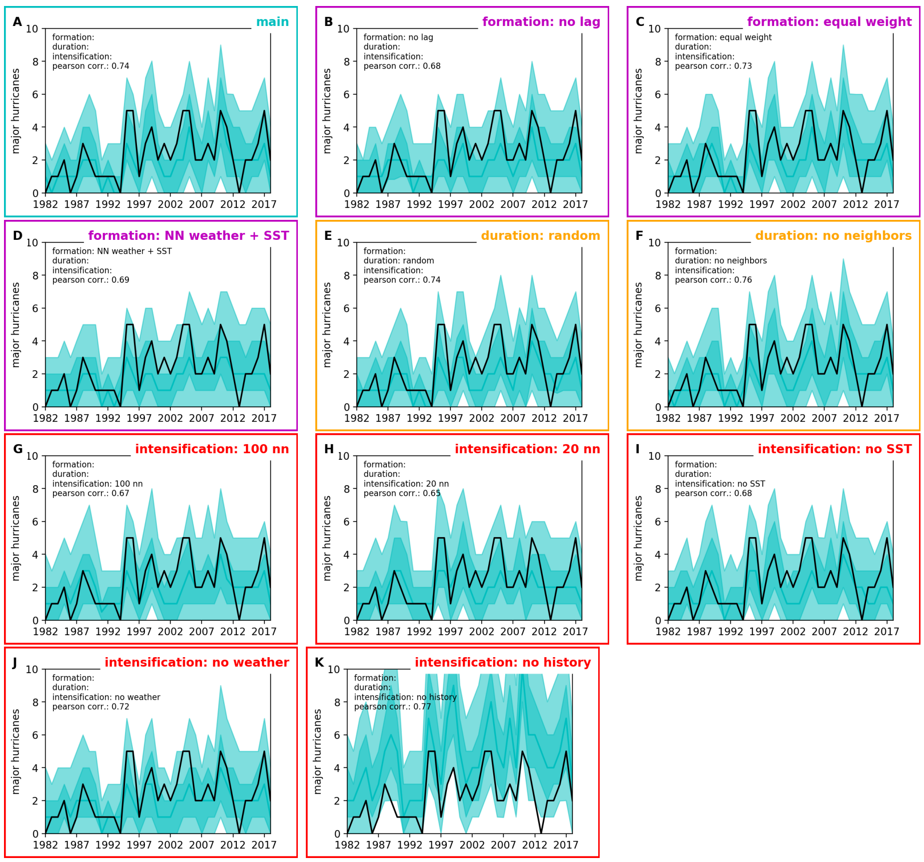

Figure S18: Simulations of seasonal major hurricanes for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

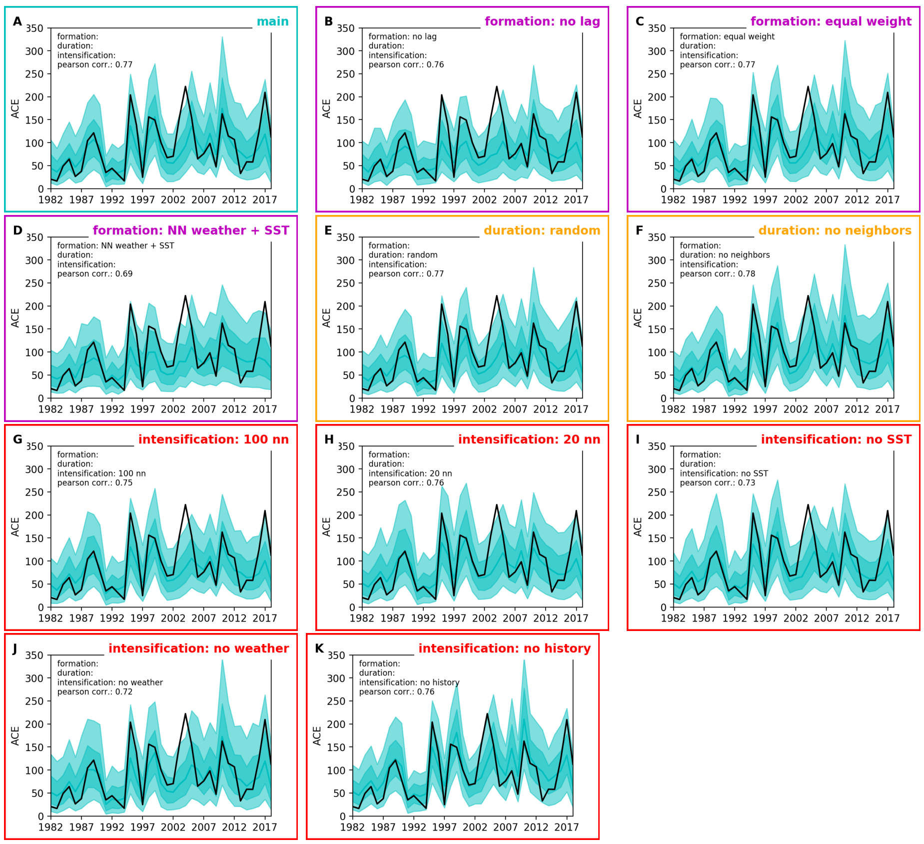

Figure S19: Simulations of seasonal ACE for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

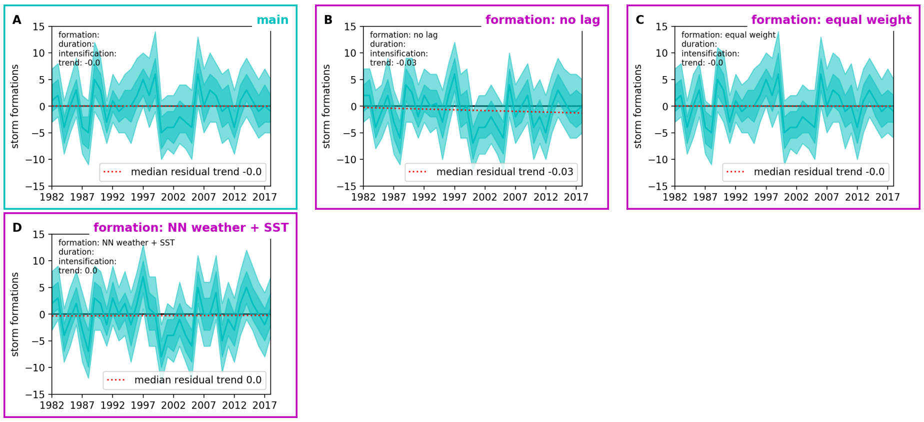

Figure S20: Residuals of simulations of seasonal storm formations for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

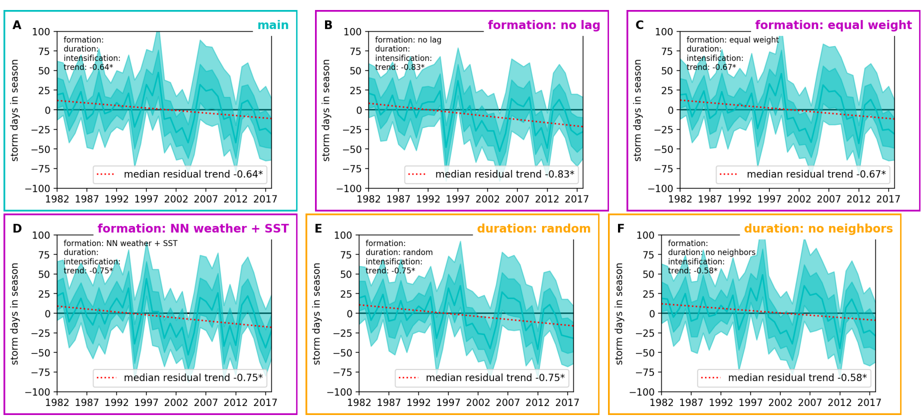

Figure S21: Residuals of simulations of seasonal storm days for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

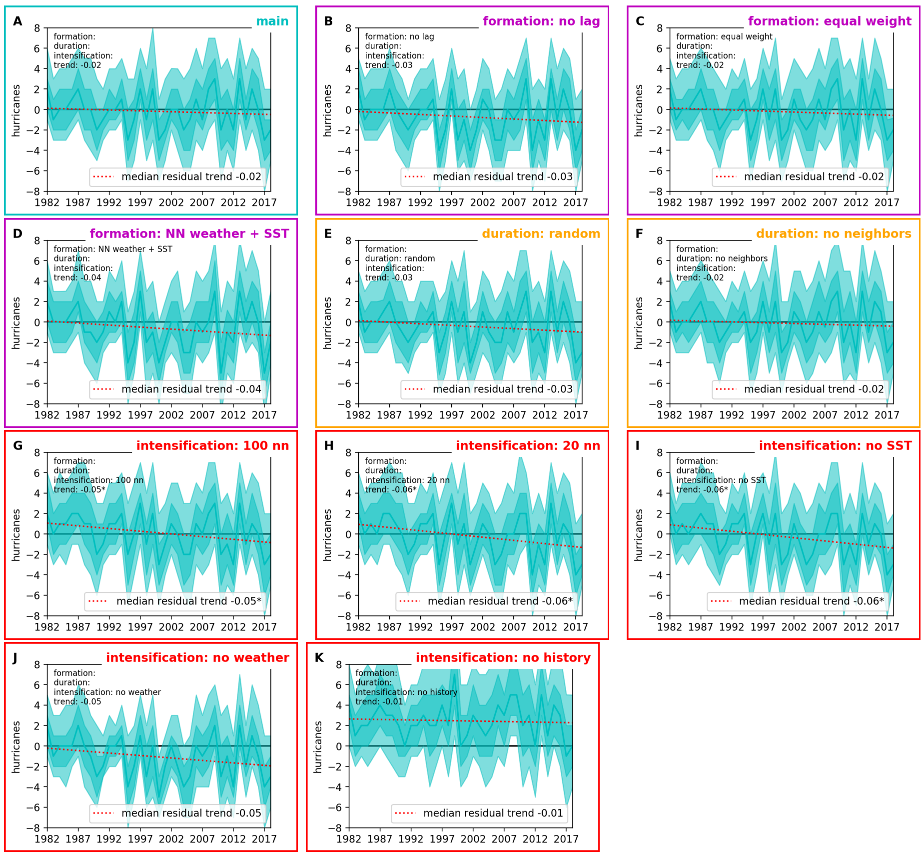

Figure S22: Residuals of simulations of seasonal hurricanes for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

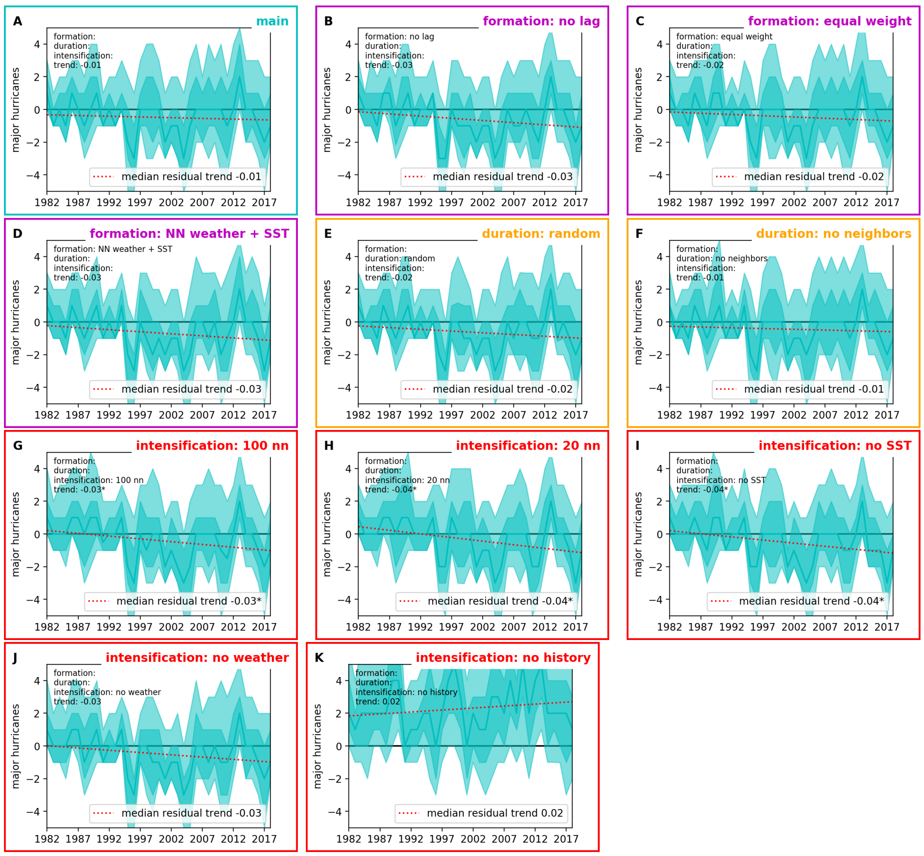

Figure S23: Residuals of simulations of seasonal major hurricanes for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

Figure S24: Residuals of simulations of seasonal ACE for different versions of the emulator. The black line shows observations. The cyan line shows the median of 1000 simulations for each year, while the lighter (darker) shading shows the 66% (%95) of simulations.

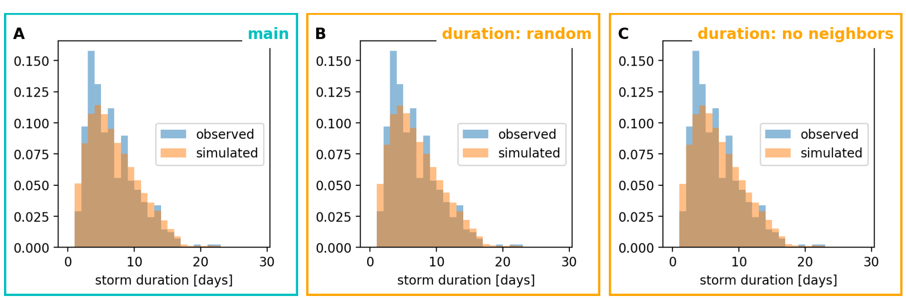

Figure S25: Histogram of storm durations for observed storms (blue) and simulated storms (orange).

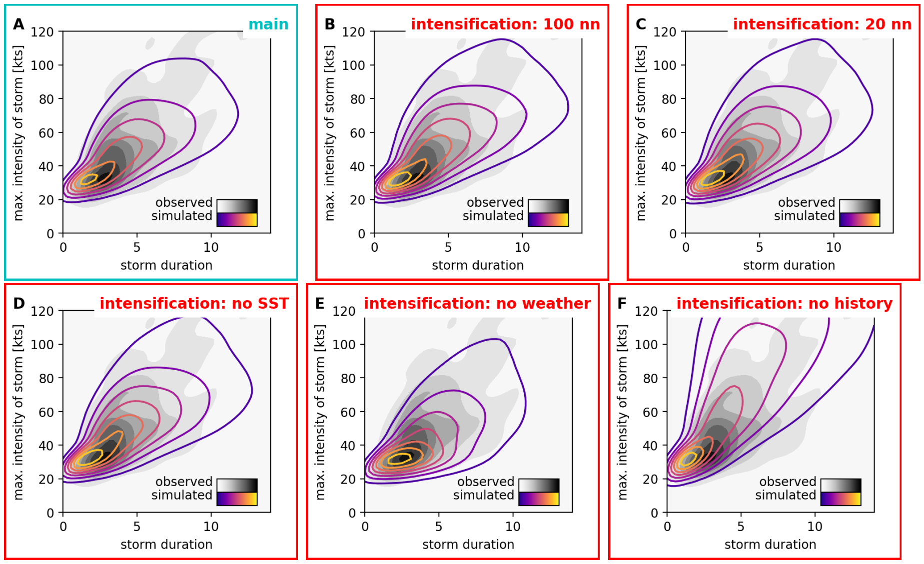

Figure S26: Maximal intensity of a storm against its duration. Gray shadings indicate kernel density estimates from observations, while colored contour-lines show the simualtions.

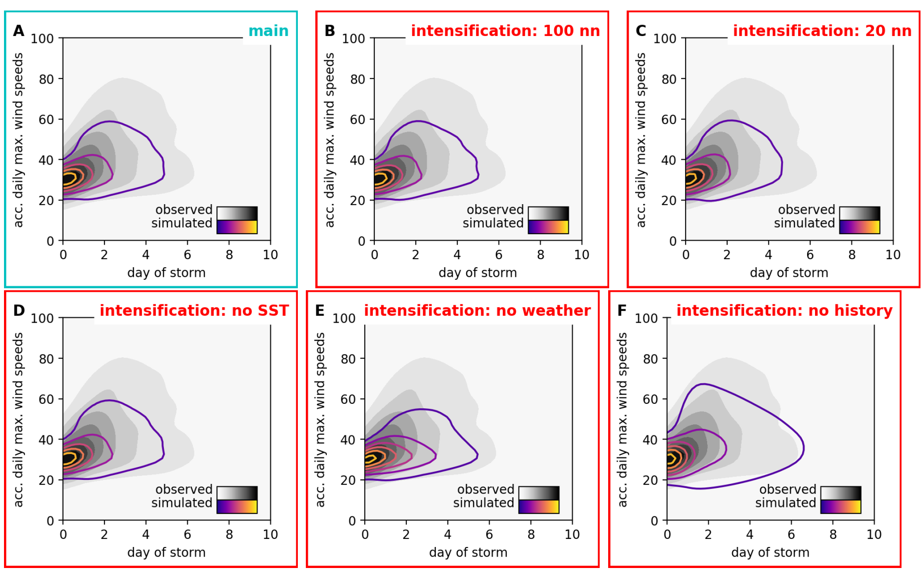

Figure S27: Daily intensity of a storm against the storm day. Gray shadings indicate kernel density estimates from observations, while colored contour-lines show the simualtions.

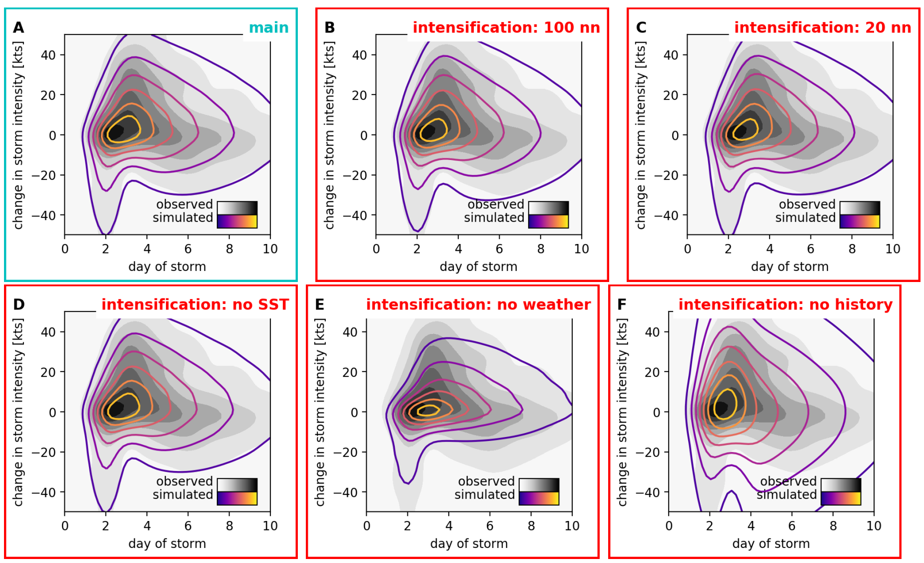

Figure S28: Daily intensity change against the storm day. Gray shadings indicate kernel density estimates from observations, while colored contour-lines show the simualtions.