Supplementary Information

Holland wind model

A Holland wind model [31] is used to estimate the axisymmetric rotational wind vectors across the domain, as a function of distance from the storm eye.

$$f = \left| \frac{4\pi sin(\phi)}{24 \cdot 60 \cdot 60} \right| \tag{1}$$

$$\beta = \frac{e\rho(V_R^2 + fV_R R_{\mathrm{MW}})}{\Delta P} \tag{2}$$

$$V(r) = \sqrt{\frac{\frac{R_{\mathrm{MW}}}{r}^{\beta} \beta \Delta P e^{-\frac{R_{\mathrm{MW}}}{r}^{\beta}}}{\rho}} - \frac{fr}{2} \tag{3}$$

where f is the Coriolis parameter, ϕ the eye latitude, β a term capturing the steepness of the radial speed profile, V_R the maximum wind speed (at the radius of maximum winds), ρ the air density, R_{MW} the radius to maximum winds, ΔP the pressure deficit of the eye below background for the region and $V(r)$ the speed at a given radius, r.

The translational motion of the storm also produces a wind component. We approximate this with the storm eye velocity vector. This vector is reduced in magnitude by α and rotated (in the direction of cyclone rotation) by β. We take radially averaged values of $\alpha = 0.56$ and $\beta = 19.2°$ as per [30].

The translational and rotational components are vector summed, giving rise to a non-axisymmetric wind field. The degree of non-axisymmetry is a function of the ratio between the storm eye speed and V_{RMW}. Where the storm eye speed is a larger proportion V_{RMW}, there will be less axisymmetry, as in the case of weaker storms [40].

Surface wind speed

To estimate surface wind speeds we use a use a spatially varying boundary layer approximation, dependent on the local surface roughness, with a power law approach as per [26].

$$V_1 = V_2 \left(\frac{z_1}{z_2} \right)^{\frac{1}{ln(\sqrt{z_1 z_2})/R_0}} \tag{4}$$

where V_1 is the velocity at height z_1, V_2 the velocity at z_2 and R_0 the surface roughness length. We took the ESA GlobCover 2009 global land cover map [41] and converted it to surface roughness values using a modified version of the values used for the Global Wind Atlas [42].

The wind footprints have been calibrated against other model data from [32]. We chose to calibrate against model rather than observation data as it is available as complete fields, allowing for many more comparisons than against point station data. Reanalysis data, such as ERA5, which combines observation and model results, could have been used, but is not available at a sufficiently high resolution to resolve the eye of a TC where wind speed varies rapidly as a function of space.

Done et al. [32] simulated IBTrACS historic storms with a gradient-level parametric wind model and then brought these winds to the surface level with a physical boundary layer model (which is too computationally intensive to practically use for thousands of events) as derived by [43]. There were 458 storms common to our work and Done's catalogue. We treated the $\frac{z_1}{z_2}$ ratio from equation 4 as a free parameter to calibrate against Done et al. This ratio was varied until the mean error (across all pixels and speeds) was minimal (with $\frac{z_1}{z_2} = 1.8$ giving $\mu_{error} = 0.2\mathrm{ms}^{-1}$).

The post-calibration pixelwise comparison of wind speed for all storms and all pixels within 200km of the storm eye is shown as Figure 7. With the population bias removed by the calibration, lower wind speeds (10-50ms^{-1}) have median errors within 2ms^{-1} of Done et al., with p_{25} and p_{75} not exceeding 5ms^{-1} absolute error. At higher wind speeds, our model increasingly underestimates wind speeds compared to Done et al., with a median error for the 60-70ms^{-1} band of approximately -5ms^{-1} and for 70-80ms^{-1}, -7ms^{-1}. For this work, electricity outage estimates will be most sensitive to wind speed errors around our transmission line failure thresholds. These are typically in the 30-40ms^{-1} bracket, with a median error of $\approx$ 1ms^{-1} and 90% of values within $\pm$7ms^{-1}.

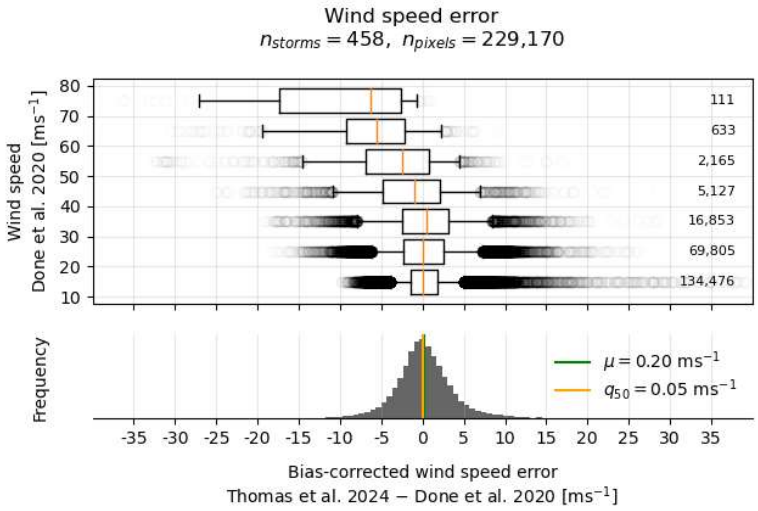

Figure 7: Post-calibration discrepancy between wind speeds generated for this work and [32]. Data drawn from 458 storms, where 229,170 pixels were within 200km of a storm eye that was over land. The lower panel shows the distribution of errors, taking Done et al. to be truth. The upper panel shows the breakdown as a function of Done et al.'s estimated wind speed. The box plots show the interquartile range and the median, with whiskers at the 5^{th} and 95^{th} percentiles and any outliers. Numbers of pixels per wind speed band are given on the right.

Calibration of wind threshold of electricity asset failure

The outage model was calibrated and validated with earth observation data from 252 and 63 storms respectively (see figure 8) that occurred around the world over the past two decades (IBTrACS 2012–2023 [13]).

The observed storm-induced electricity outage was estimated using NASA's Black Marble global daily nocturnal visible light measurements at 15 arc second (VNP46A2) [44]. For each storm, the before-storm baseline was computed as the average of the measurements from 3 months before the storm; and the after-storm NTL was estimated with a 10-day composite image (i.e., pixels of the image were filled with the first available persistent night-time lights, up to 10 days after the storm). For the calibration storms we computed network disruption results for eleven wind speed damage thresholds, v_{th} (20 to 45ms^{-1} in 2.5ms^{-1}) increments). The outage model performance was evaluated by comparing the relative changes in before and after storm night-time lights and the modelled power supply faction, for each function of the wind speed damage thresholds, using Franks' index [45].

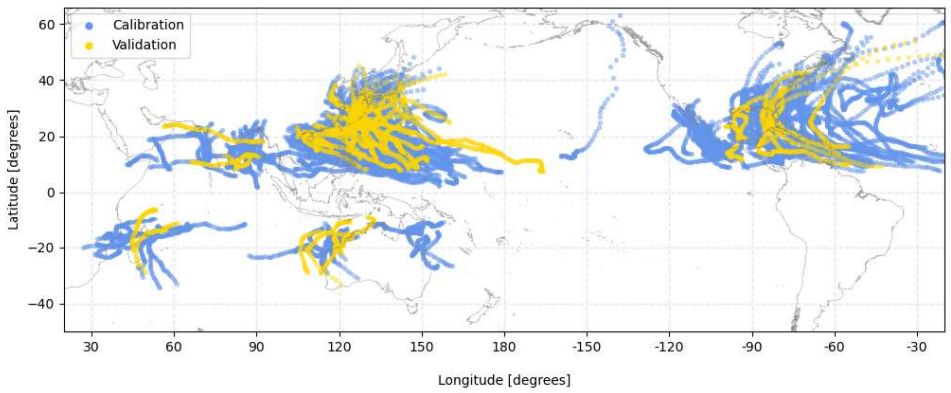

Figure 8: IBTrACS historic tracks as used for calibration and validation of wind speed damage thresholds.

$$F_{th} = 1 - \frac{|A_{mod,th} - A_{obs}|}{A_{obs}} \tag{5}$$

where $A_{mod,th}$ represents the number of electricity consuming sinks predicted as impacted for a given wind speed damage threshold and A_{obs} the number of sinks observed to be impacted from the NTL data. F_{th} is Frank's index for a given wind speed damage threshold and ranges from $-\infty$ to 1, with a value closer to 1 indicating better model performance. This is shown in Figure 6 e), where the medium damage threshold attains the highest Frank's index as it reproduces the correct number of disconnected sinks.

We define a simple ratio, r, to compare electricity consuming sinks before and after storms and identify observed and modelled events. The definition of an observed event is when the post-storm NTL intensity, NL_{post} is sufficiently below the the pre-storm value, NL_{pre}, that is, when $\mathrm{NL}_{post} < r\mathrm{NL}_{pre}$. Similarly, for the modelled results, an outage is defined as when a sink's connected power generating supply, s_{pre} is sufficiently less than the pre-storm connected power generating supply, s_{post}, i.e. $s_{post} < rs_{pre}$. In both observed and modelled cases, we take $r = 0.9$.

We have performed a sensitivity analysis of s in the range 0.7-0.9, and found that our results are not sensitive to the values selected (i.e., the best wind thresholds identified are the same under different values of s). Where sufficient calibration data was present for a given country, we use country-specific thresholds. These are 25ms^{-1} (Mexico, Madagascar, St. Lucia, Guatemala, Guam, India, Thailand, Laos, Australia), 27.5ms^{-1} (Dominican Republic, Vietnam), 32.5ms^{-1} (China, Jamaica, Canada, Cuba), 35ms^{-1}(USA, Korea, Philippines, Hong Kong, Japan) and 42.5^{-1} (Taiwan). These calibrated damage thresholds capture some of the inherent resilience (or otherwise) of the country networks, but they also act to compensate for modelling errors. For example, our wind model does not take account of orographic effects, so the shelter that Taiwanese mountains provide for the populated western seaboard from Pacific tropical cyclones is not modelled well. Our relatively high damage threshold of 42.5ms^{-1} for Taiwan counteracts this. In the case where we do not have enough historic storm data to calibrate a country, we fall back to a threshold calculated across a World Bank income classification – for high-income countries a threshold of 32.5ms^{-1} and 30ms^{-1} for other income groups.

Validation using the 20% of storms not employed in calibration yielded an accuracy of $0.79 \leq F \leq 0.88$.